Matrix factorizations and low rank approximation

The first section of the course provides a quick review of basic concepts from linear algebra that we will use frequently. Note that the pace is fast here, and assumes that you have seen these concepts in prior course-work. If not, then additional reading on the side is strongly recommended!

1. Notation, etc

1.1. **Norms.** Let $\mathbf{x} = [x_1, x_2, \ldots, x_n]$ denote a vector in $\mathbb{R}^n$ or $\mathbb{C}^n$. Our default norm for vectors is the Euclidean norm

$$\|\mathbf{x}\| = \left(\sum_{j=1}^{n} |x_j|^2\right)^{1/2}.$$

We will at times also use ℓ^p norms

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

Let $\mathbf{A}$ denote an $m \times n$ matrix. For the most part, we allow $\mathbf{A}$ to have complex entries. We define the *spectral norm* of $\mathbf{A}$ via

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

We define the *Frobenius norm* of $\mathbf{A}$ via

$$\|\mathbf{A}\|_{\mathrm{F}} = \left(\sum_{i=1}^{m}\sum_{j=1}^{n} |\mathbf{A}(i,j)|^2\right)^{1/2}.$$

Observe that

$$\|\mathbf{A}\| \leq \|\mathbf{A}\|_{\mathrm{F}} \leq \sqrt{\min(m,n)}\, \|\mathbf{A}\|.$$

1.2. **Transpose and adjoint.** Given an $m \times n$ matrix real $\mathbf{A}$, the *transpose* $\mathbf{A}^{\mathrm{t}}$ is the $n \times m$ matrix $\mathbf{B}$ with entries

$$\mathbf{B}(i,j) = \mathbf{A}(j,i).$$

If $\mathbf{A}$ is *complex*, then the natural analogue of the transpose is the *adjoint*, which is the complex conjugate of the transpose

$$\mathbf{A}^* = \overline{\mathbf{A}^{\mathrm{t}}}.$$

1.3. **Subspaces.** Let $\mathbf{A}$ be an $m \times n$ matrix.

- The *row space* of $\mathbf{A}$ is denoted $\mathrm{row}(\mathbf{A})$ and is defined as the subspace of $\mathbb{R}^n$ spanned by the rows of $\mathbf{A}$.
- The *column space* of $\mathbf{A}$ is denoted $\mathrm{col}(\mathbf{A})$ and is defined as the subspace of $\mathbb{R}^m$ spanned by the columns of $\mathbf{A}$. The column space is the same as the *range* or $\mathbf{A}$, so $\mathrm{col}(\mathbf{A}) = \mathrm{ran}(\mathbf{A})$.
- The *nullspace* or *kernel* of $\mathbf{A}$ is the subspace $\ker(\mathbf{A}) = \mathrm{null}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\}$.

1.4. **Special classes of matrices.** We use the following terminology to classify matrices:

- An $m \times n$ matrix $\mathbf{A}$ is *orthonormal* if its columns form an orthonormal basis, i.e. $\mathbf{A}^*\mathbf{A} = \mathbf{I}$.
- An $n \times n$ matrix $\mathbf{A}$ is *normal* if $\mathbf{A}\mathbf{A}^* = \mathbf{A}^*\mathbf{A}$.
- An $n \times n$ real matrix $\mathbf{A}$ is *symmetric* if $\mathbf{A}^{\mathrm{t}} = \mathbf{A}$.
- An $n \times n$ matrix $\mathbf{A}$ is *self-adjoint* if $\mathbf{A}^* = \mathbf{A}$.
- An $n \times n$ matrix $\mathbf{A}$ is *skew-adjoint* if $\mathbf{A}^* = -\mathbf{A}$.
- An $n \times n$ matrix $\mathbf{A}$ is *unitary* if it is invertible and $\mathbf{A}^* = \mathbf{A}^{-1}$.

2. Low rank approximation

2.1. **Exact rank deficiency.** Let $\mathbf{A}$ be an $m \times n$ matrix. Let k denote an integer between 1 and $\min(m, n)$. Then the following conditions are equivalent:

- The columns of $\mathbf{A}$ span a subspace of $\mathbb{R}^m$ of dimension k.
- The rows of $\mathbf{A}$ span a subspace of $\mathbb{R}^n$ of dimension k.
- The nullspace of $\mathbf{A}$ has dimension $n - k$.
- The nullspace of $\mathbf{A}^*$ has dimension $m - k$.

If $\mathbf{A}$ satisfies any of these criteria, then we say that $\mathbf{A}$ *has rank* k. When $\mathbf{A}$ has rank k, it is possible to find matrices $\mathbf{E}$ and $\mathbf{F}$ such that

$$\underset{m \times n}{\mathbf{A}} \approx \underset{m \times k}{\mathbf{E}} \underset{k \times n}{\mathbf{F}}.$$

Having access to such factors $\mathbf{E}$ and $\mathbf{F}$ can be very helpful:

- Storing $\mathbf{A}$ requires mn words of storage.
  Storing $\mathbf{E}$ and $\mathbf{F}$ requires $km + kn$ words of storage.
- Given a vector $\mathbf{x}$, computing $\mathbf{Ax}$ requires mn flops.
  Given a vector $\mathbf{x}$, computing $\mathbf{Ax} = \mathbf{E}(\mathbf{Fx})$ requires $km + kn$ flops.
- The factors $\mathbf{E}$ and $\mathbf{F}$ are often useful for *data interpretation.*

In practice, we often impose conditions on the factors. For instance, in the well known QR decomposition, the columns of $\mathbf{E}$ are orthonormal, and $\mathbf{F}$ is upper triangular (up to permutations of the columns).

2.2. **Approximate rank deficiency.** The condition that $\mathbf{A}$ has *precisely* rank k is of high theoretical interest, but is not realistic in practical computations. Frequently, the numbers we use have been measured by some device with finite precision, or they may have been computed via a simulation with some approximation errors (e.g. by solving a PDE numerically). In any case, we almost always work with data that is stored in some finite precision format (typically about 10^{-15}). For all these reasons, it is very useful to define the concept of *approximate rank.* In this course, we will typically use the following definition:

Definition 2.1. Let $\mathbf{A}$ be an $m \times n$ matrix, and let ε be a positive real number. We then say define the ε-rank k of $\mathbf{A}$ as the unique integer k such that both the following two conditions holds:

(a) There exists a matrix $\mathbf{B}$ of precise rank k such that $\|\mathbf{A} - \mathbf{B}\| \leq \varepsilon$.
(b) There does not exist any matrix $\mathbf{B}$ of rank less than k such that (a) holds

The term ε-rank is sometimes used slightly loosely without enforcing condition (b). In other words, we sometimes say that $\mathbf{A}$ has ε-rank k if

$$\inf\{\|\mathbf{A} - \mathbf{B}\| : \mathbf{B} \text{ has rank } k\},$$

without worrying about whether the rank could actually be smaller.

3. The eigenvalue decomposition

Let $\mathbf{A}$ be an $n \times n$ matrix (it must be *square* for eigenvalues and eigenvectors to exist). We then say that λ is an eigenvalue and $\mathbf{v}$ is an eigenvector of $\mathbf{A}$ if $\mathbf{v} \neq \mathbf{0}$ and

$$\mathbf{Av} = \lambda \mathbf{v}.$$

Theorem 1. *Let $\mathbf{A}$ be an $n \times n$ normal matrix (so that $\mathbf{AA}^* = \mathbf{A}^*\mathbf{A}$). Then there exist a unitary matrix $\mathbf{V}$ and a diagonal matrix $\mathbf{D}$ such that*

$$\mathbf{A} = \mathbf{VDV}^*. \tag{1}$$

The equation (1) can alternatively be written

$$\mathbf{A} = \sum_{j=1}^{n} \lambda_j \, \mathbf{v}_j \, \mathbf{v}_j^*,$$

where $\{\lambda_j, \, \mathbf{v}_j\}$ are the *eigenpairs* of $\mathbf{A}$, and

$$\mathbf{V} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_p \end{bmatrix},$$

$$\mathbf{D} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \lambda_p \end{bmatrix}.$$

In other words, at the columns of $\mathbf{V}$ are the eigenvectors of $\mathbf{A}$. These eigenvectors form an orthonormal basis for $\mathbb{C}^n$. In the basis $\{\mathbf{v}_j\}_{j=1}^n$, the matrix $\mathbf{A}$ is diagonal.

Recall that *self-adjoint*, *skew-adjoint*, and *unitary* matrices are special cases of *normal* matrices, so these classes all allow spectral decompositions. It is easy to verify that:

$\mathbf{A}$ is self-adjoint	$\Leftrightarrow$	$\mathbf{A}^* = \mathbf{A}$	$\Leftrightarrow$	Every eigenvalue is real.
$\mathbf{A}$ is skew-adjoint	$\Leftrightarrow$	$\mathbf{A}^* = -\mathbf{A}$	$\Leftrightarrow$	Every eigenvalue is imaginary.
$\mathbf{A}$ is unitary	$\Leftrightarrow$	$\mathbf{A}^* = \mathbf{A}^{-1}$	$\Leftrightarrow$	Every eigenvalue satisfies $\lvert\lambda_j\rvert = 1$.

Note that even a matrix whose entries are all real may have complex eigenvalues and eigenvectors.

What about non-normal matrices? Every square matrix has at least one (possibly complex) eigenvalue and one eigenvector. But if $\mathbf{A}$ is not normal, then there is no orthonormal basis consisting of all eigenvectors. While eigenvalue decompositions of non-normal matrices are still very useful for certain applications, the lack of an ON-basis consisting of eigenvectors makes the *singular value decomposition* a much better tool for low-rank approximation in this case.