Week 2: Monday, Aug 27

Matrix norms

In the last lecture, we discussed norms and inner products on vector spaces. Spaces of linear maps (or matrices) can also be treated as vector spaces, and the same definition of norms applies. In general, though, we would like to consider norms on spaces of linear maps that are in some way compatible with the norms on the spaces they map between.

If A maps between two normed vector spaces $\mathcal{V}$ and $\mathcal{W}$, the *induced norm* on A is

$$\|A\|_{\mathcal{V},\mathcal{W}} = \sup_{v \neq 0} \frac{\|Av\|_{\mathcal{W}}}{\|v\|_{\mathcal{V}}}.$$

Because norms are homogeneous with respect to scaling, we also have

$$\|A\|_{\mathcal{V},\mathcal{W}} = \sup_{\|v\|_{\mathcal{V}}=1} \|Av\|_{\mathcal{W}}.$$

Note that when $\mathcal{V}$ is finite-dimensional (as it always is in this class), the unit ball $\{v \in \mathcal{V} : \|v\| = 1\}$ is compact, and $\|Av\|$ is a continuous function of v, so the supremum is actually attained.

These operator norms are indeed norms on the space $\mathcal{L}(\mathcal{V}, \mathcal{W})$ of bounded linear maps between $\mathcal{V}$ and $\mathcal{W}$ (or norms on vector spaces of matrices, if you prefer). Such norms have a number of nice properties, not the least of which are the submultiplicative properties

$$\begin{aligned} \|Av\| &\leq \|A\|\|v\| \\ \|AB\| &\leq \|A\|\|B\|. \end{aligned}$$

The first property ($\|Av\| \leq \|A\|\|v\|$) is clear from the definition of the vector norm. The second property is almost as easy to prove:

$$\|AB\| = \max_{\|v\|=1} \|ABv\| \leq \max_{\|v\|=1} \|A\|\|Bv\| = \|A\|\|B\|.$$

The matrix norms induced when $\mathcal{V}$ and $\mathcal{W}$ are supplied with a 1-norm, 2-norm, or ∞-norm are simply called the matrix 1-norm, 2-norm, and ∞-norm.

The matrix 1-norm and ∞-norm are given by

$$\|A\|_1 = \max_j \sum_i |A_{ij}|$$
$$\|A\|_\infty = \max_i \sum_j |A_{ij}|.$$

These norms are nice because they are easy to compute. Also easy to compute (though it's not an induced operator norm) is the *Frobenius* norm

$$\|A\|_F = \sqrt{\operatorname{tr}(A^*A)} = \sqrt{\sum_{i,j} |A_{ij}|^2}.$$

The Frobenius norm is not an operator norm, but it does satisfy the submultiplicative property (i.e. it is *consistent* with the vector 2-norm).

The 2-norm

The matrix 2-norm is very useful, but it is also not so straightforward to compute. However, it has an interesting characterization. If A is a real matrix, then we have

$$\begin{aligned} \|A\|_2^2 &= \left(\max_{\|v\|_2=1} \|Av\| \right)^2 \\ &= \max_{\|v\|_2^2=1} \|Av\|^2 \\ &= \max_{v^T v=1} v^T A^T A v. \end{aligned}$$

This is a constrained optimization problem, to which we will apply the method of Lagrange multipliers: that is, we seek critical points for the functional

$$L(v, \mu) = v^T A^T A v - \mu(v^T v - 1).$$

Differentiate in an arbitrary direction $(\delta v, \delta \mu)$ to find

$$\begin{aligned} 2\delta v^T (A^T A v - \mu v) &= 0, \\ \delta \mu (v^T v - 1) &= 0. \end{aligned}$$

Therefore, the stationary points satisfy the eigenvalue problem

$$A^T A v = \mu v.$$

The eigenvalues of $A^T A$ are non-negative (why?), so we will call them σ_i^2. The positive values σ_i are called the *singular values* of A, and the largest of these singular values is $\|A\|_2$. We will return to the idea of singular values, and the properties we can infer from them, in the not-too-distant future.

Error measures and norms

One reason we care about norms is because they give us a concise way of talking about the sizes of errors and perturbations. Before giving an example, let us set the stage with some basic definitions.

First, suppose $\hat{\alpha}$ and α are scalars, with $\hat{\alpha} \approx \alpha$. The *absolute error* in $\hat{\alpha}$ as an approximation to α is simply $|\hat{\alpha} - \alpha|$. In some cases, we will refer to this simply as "the error" (possibly leaving off the absolute value). Unfortunately, absolute errors are difficult to interpret out of context. For example, suppose that I have a measurement with an absolute error of one meter. If this is a measurement of my height, a meter is a very large error; if it is the distance from the earth to the sun, the error is tiny. For this reason, we most often use the dimensionless *relative error*, $|\hat{\alpha} - \alpha|/|\alpha|$.

For vectors, we can likewise talk about the *normwise absolute error* $\|\hat{x}-x\|$ or the *normwise relative error* $\|\hat{x} - x\|/\|x\|$. Less frequently, we will discuss the *componentwise absolute errors* $\hat{x}_i - x_i$ and the *componentwise relative errors* $|\hat{x}_i - x_i|/|x_i|$. Note that the maximum of the componentwise absolute errors is simply $\|\hat{x} - x\|_\infty$, and the maximum of the componentwise relative errors is $\|\operatorname{diag}(x)^{-1}(\hat{x} - x)\|_\infty$

Introducing the condition number

Now, suppose I want to compute $y = Ax$, but because of a small error in A (due to measurement errors or roundoff effects), I instead compute $\hat{y} = (A + E)x$ where E is "small." Of course, the expression for the *absolute* error is trivial:

$$\|\hat{y} - y\| = \|Ex\|.$$

But I usually care more about the *relative error*.

$$\frac{\|\hat{y} - y\|}{\|y\|} = \frac{\|Ex\|}{\|y\|}.$$

If we assume that A is invertible and that we are using consistent norms (which we will usually assume), then

$$\|Ex\| = \|EA^{-1}y\| \leq \|E\|\|A^{-1}\|\|y\|,$$

which gives us

$$\frac{\|\hat{y} - y\|}{\|y\|} \leq \|A\|\|A^{-1}\|\frac{\|E\|}{\|A\|} = \kappa(A)\frac{\|E\|}{\|A\|}.$$

That is, the relative error in the output is the relative error in the input multiplied by the *condition number* $\kappa(A) = \|A\|\|A^{-1}\|$. This concept of a condition number as a relation between two relative errors will be a recurring theme in the class.