Linear Algebra

We shall turn our attention to solving linear systems of equations

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

where $\mathbf{A} \in \mathbb{R}^{m\times n},\ \boldsymbol{x} \in \mathbb{R}^n,$ and $\boldsymbol{b} \in \mathbb{R}^m$.

We already saw examples of methods that required the solution of a linear system as part of the overall algorithm, e.g. the Vandermonde system for interpolation, which was a square system ($m = n$).

Another category of methods that leads to *rectangular* systems with $m > n$ is *least square methods*. They answer questions of the form:

- What is the *best* n-order polynomial we can use to *approximate* (not interpolate) $m + 1$ data points (where $m > n$)
- More generally, find the solution that *most closely* satisfies m equations, in the presence of n ($n < m$) unknowns.

Lecture of: 19 Mar 2013

All these algorithms need to be conscious about *error* and there are at least 3 sources where error stems from:

- Some algorithms are "imperfect" in the sense that they require several iterations to generate a good quality approximation. Thus, intermediate results are subject to *error*.
- Sometimes, it is not possible to fine an "ideal" solution, e.g. because we have more equations than unknowns. In this case, not all equations will be satisfied exactly, and we need a notion of the "error" incurred in not satisfying certain equations fully.
- Inputs to an algorithm are often corrupted by noise, roundoff error, etc. For example, instead of solving an "intended" system $\mathbf{A}\boldsymbol{x} = \boldsymbol{b}$ we may be solving $\mathbf{A}^*\boldsymbol{x} = \boldsymbol{b}^*$ where the entries $\mathbf{A}^*$ and $\boldsymbol{b}^*$ have been subject to noise and inaccuracy. It is important to know how these translate to errors in determining $\boldsymbol{x}$.

Vector Norms

Norms are valuable tools in arguing about the extent and magnitude of error. We will introduce some concepts that we will use broadly later on.

Definition: A vector norm is a function from $\mathbb{R}^n$ to $\mathbb{R}$, with a certain number of properties. If $\boldsymbol{x} \in \mathbb{R}^n$, we symbolize its norm by $||\boldsymbol{x}||$. The defining properties of a norm are :

(i) $||\boldsymbol{x}|| \geq 0$ for all $\boldsymbol{x} \in \mathbb{R}^n$ and also $||\boldsymbol{x}|| = 0$ iff $\boldsymbol{x} = \boldsymbol{0}$

(ii) $||\alpha \boldsymbol{x}|| = |\alpha| \cdot ||\boldsymbol{x}||$ for all $\alpha \in \mathbb{R}$, $\boldsymbol{x} \in \mathbb{R}^n$

(iii) $||\boldsymbol{x} + \boldsymbol{y}|| \leq ||\boldsymbol{x}|| + ||\boldsymbol{y}||$ for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ (triangle inequality)

Note that the properties above do not determine a unique form of a *norm* function, in fact many different valid norms exist. Typically we will use subscripts ($|| \cdot ||_a, || \cdot ||_b$) to denote different types of norms.

So, *why* do we need to introduce the concept of a vector norm? The reason has to do with the following fact:

When dealing with the solution of a nonlinear equation $f(x) = 0$, the error $e = x_{\text{approx}} - x_{\text{exact}}$ is a single number, thus the absolute value $|e|$ gives us a good idea of the "extent" of the error.

When solving a system of linear equations $A\boldsymbol{x} = \boldsymbol{b}$, the exact solution $\boldsymbol{x}_{\text{exact}}$ as well as any approximation $\boldsymbol{x}_{\text{approx}}$ are vectors, and the error:

$$\boldsymbol{e} = \boldsymbol{x}_{\text{approx}} - \boldsymbol{x}_{\text{exact}}$$

is a vector, too. It is not straightforward to assess the "magnitude" of such a vector-valued error. e.g. Consider $\boldsymbol{e}_1, \boldsymbol{e}_2 \in \mathbb{R}^{1000}$, and

$$\boldsymbol{e}_1 = \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ \vdots \\ 0.1 \end{pmatrix} \quad \boldsymbol{e}_2 = \begin{pmatrix} 100 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

Which one is worse? $\boldsymbol{e}_1$ has a modest amount of error, distributed over all components. In $\boldsymbol{e}_2$, all but one component are exact, but one of them has a huge discrepancy.

Exactly how we quantify and assess the extent of error is application-dependent. Vector norms are alternative ways to measure this magnitude and different norms would be appropriate for different tasks. There are, however, some specific definitions of vector norms, which are commonly used. All of them can be proven to satisfy the norm properties, and their definitions are:

(In all definitions below, $\boldsymbol{x} = (x_1, x_2, \cdots, x_n)$)

1. The L_1-norm (or 1-norm)
$$||\boldsymbol{x}||_1 = \sum_{i=1}^{n} |x_i|$$
2. The L_2-norm (or 2-norm, or Euclidean norm)
$$||\boldsymbol{x}||_2 = \sqrt{\sum_{i=1}^{n} x_i^2}$$
3. The infinity norm (or max-norm)
$$||\boldsymbol{x}||_\infty = \max_{1\le i\le n} |x_i|$$
4. (Less common) L_p norm
$$||\boldsymbol{x}||_p = \left(\sum_{i=1}^{n} |x_i|^p\right)^{1/p}$$

It is relatively easy to show that these satisfy the defining properties of a norm e.g. for $||\cdot||_1$:

- $||\boldsymbol{x}|| = \sum_{i=1}^{n} |x_i| \ge 0$
- if $\boldsymbol{x} = 0$, then $||\boldsymbol{x}|| = 0$
  if $||\boldsymbol{x}|| = 0 \Rightarrow \sum_{i=1}^{n} |x_i| = 0 \Rightarrow |x_i| = 0 \ \forall i \Rightarrow \boldsymbol{x} = 0$
- $||\alpha\boldsymbol{x}|| = \sum_{i=1}^{n} |\alpha x_i| = |\alpha| \sum_{i=1}^{n} |x_i| = |\alpha| \, ||\boldsymbol{x}||_1$

Similar proofs can be given for $||\cdot||_\infty$ (just as easy), $||\cdot||_2$ (a bit more difficult) and $||\cdot||_p$ (rather complicated).

Matrix Norms

We can actually define norms for (square) matrices, as well.

Definition: A matrix norm is a function from $\mathbb{R}^{n\times n}$ to $\mathbb{R}$ which satisfies a certain number of properties. If $\mathbf{A} \in \mathbb{R}^{n\times n}$, we symbolize its norm by $||\mathbf{A}||$. The defining properties of a matrix norm are :

(i) $||\mathbf{M}|| \geq 0$ for all $\mathbf{M} \in \mathbb{R}^{n \times n}$, and also $||\mathbf{M}|| = 0$ iff $\mathbf{M} = \mathbf{0}$

(ii) $||\alpha\mathbf{M}|| = |\alpha| \cdot ||\mathbf{M}||$ for all $\alpha \in \mathbb{R}$

(iii) $||\mathbf{M} + \mathbf{N}|| \leq ||\mathbf{M}|| + ||\mathbf{N}||$

(iv) $||\mathbf{M} \cdot \mathbf{N}|| \leq ||\mathbf{M}|| \cdot ||\mathbf{N}||$

Property (iv) is the one that has slightly different flavor than the properties of vector norms. Although more types of matrix norms do exist, one common category is that of matrix norms *induced* by vector norms.

Definition: If $|| \cdot ||_*$ is a valid vector norm, its *induced* matrix norm is defined as

$$||\mathbf{M}||_* = \max_{\substack{\boldsymbol{x} \in \mathbb{R}^n \\ \boldsymbol{x} \neq 0}} \left\{ \frac{||M\boldsymbol{x}||_*}{||\boldsymbol{x}||_*} \right\}$$

or equivalently:

$$||\mathbf{M}||_* = \max_{\substack{\boldsymbol{x} \in \mathbb{R}^n \\ ||\boldsymbol{x}|| = 1}} \{||M\boldsymbol{x}||_*\}$$

Note, again, that *not all* valid matrix norms are induced by vector norms. One notable example is the very commonly used *Frobenius norm*:

$$||\mathbf{M}||_F = \sqrt{\sum_{i,j=1}^{n} M_{ij}^2}$$

We can easily show though that induced norms satisfy properties (i) through (iv). (i)-(iii) are rather trivial, e.g,:

$$\begin{aligned} ||\mathbf{M} + \mathbf{N}|| &= \max_{\boldsymbol{x} \neq 0} \frac{||(\mathbf{M} + \mathbf{N})\boldsymbol{x}||}{||\boldsymbol{x}||} \leq \max_{\boldsymbol{x} \neq 0} \frac{||\mathbf{M}\boldsymbol{x}|| + ||\mathbf{N}\boldsymbol{x}||}{||\boldsymbol{x}||} \\ &= \max_{\boldsymbol{x} \neq 0} \frac{||\mathbf{M}\boldsymbol{x}||}{\boldsymbol{x}} + \max_{\boldsymbol{x} \neq 0} \frac{||\mathbf{N}\boldsymbol{x}||}{\boldsymbol{x}} = ||\mathbf{M}|| + ||\mathbf{N}|| \end{aligned}$$

Property (iv) is slightly trickier to show. First a lemma:

Lemma: If $|| \cdot ||$ is a matrix norm induced by a vector norm $|| \cdot ||$, then:

$$||\mathbf{A}\boldsymbol{x}|| \leq ||\mathbf{A}|| \cdot ||\boldsymbol{x}||$$

Proof: Since $||\mathbf{A}|| = max_{\boldsymbol{x}\neq 0}(||\mathbf{A}\boldsymbol{x}||/||\boldsymbol{x}||)$, we have that for an arbitrary $\boldsymbol{y} \in \mathbb{R}^n$(with $\boldsymbol{y} \neq 0$) :

$$||\mathbf{A}|| = \max_{\boldsymbol{x}\neq 0} \frac{||\mathbf{A}\boldsymbol{x}||}{||\boldsymbol{x}||} \geq \frac{||\mathbf{A}\boldsymbol{y}||}{||\boldsymbol{y}||} \Rightarrow$$

$$||\mathbf{A}\boldsymbol{y}|| \leq ||\mathbf{A}|| \cdot ||\boldsymbol{y}||$$

This holds for $\boldsymbol{y} \neq 0$, but we can see it is also true for $\boldsymbol{y} = 0$.
We can now show property (iv):

$$\begin{aligned} ||\mathbf{MN}|| &= \max_{||\boldsymbol{x}||=1} ||\mathbf{MN}\boldsymbol{x}|| \leq \max_{||\boldsymbol{x}||=1} ||\mathbf{M}|| \cdot ||\mathbf{N}\boldsymbol{x}|| \\ &= ||\mathbf{M}|| \cdot \max_{||\boldsymbol{x}||=1} ||\mathbf{N}\boldsymbol{x}|| = ||\mathbf{M}|| \cdot ||\mathbf{N}|| \Rightarrow \end{aligned}$$

$$||\mathbf{MN}|| \leq ||\mathbf{M}|| \cdot ||\mathbf{N}||$$

Although the definition of an induced norm allowed us to prove certain properties it does not necessarily provide a convenient formula for evaluating the matrix norm.

Fortunately, such formulas do exist for the L_1 and L_∞ induced matrix norms. Given here without proof:

$$\begin{aligned} ||\mathbf{A}||_1 &= \max_j \sum_{i=1}^{n} |A_{ij}| \text{ max. absolute column sum} \\ ||\mathbf{A}||_\infty &= \max_i \sum_{i=1}^{n} |A_{ij}| \text{ max. absolute row sum} \end{aligned}$$

The formula for $||\cdot||_2$ is much more involved; we will settle with the general definition instead of a closed-form formula in this case.

Lecture of: 21 Mar 2013

Useful Properties of Matrix and Vector Norms We previously saw that

$$||\mathbf{A}\boldsymbol{x}|| \leq ||\mathbf{A}|| \cdot ||\boldsymbol{x}|| \tag{11}$$

for any matrix $\mathbf{A}$ and any vector $\boldsymbol{x}$ of dimensions $m \times m$ and $m \times 1$, respectively).

Note that, when writing an expression such as 11, the matrix norm $||\mathbf{A}||$ is understood to be the inferred norm from the vector norm used in $||\mathbf{A}\boldsymbol{x}||$ and $||\boldsymbol{x}||$. Thus,

$$||\mathbf{A}\boldsymbol{x}||_1 \leq ||\mathbf{A}||_1 \cdot ||\boldsymbol{x}||_1$$

and

$$||\mathbf{A}\boldsymbol{x}||_\infty \leq ||\mathbf{A}||_\infty \cdot ||\boldsymbol{x}||_\infty$$

are both valid but we can not mix and match, e.g.:

$$\cancel{||\mathbf{A}\boldsymbol{x}||_\infty \leq ||\mathbf{A}||_2 \cdot ||\boldsymbol{x}||_1} \qquad \textbf{NOT CORRECT}$$