CS412: Lecture #15

Mridul Aanjaneya

March 12, 2015

We shall turn our attention to solving linear equations

$$Ax = b$$

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

All these algorithms need to be conscious about *error*, and there are at least three sources for it.

- 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 find 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 computed by noise, roundoff error, etc. For example, instead of solving an "intended" system $AX = b$ we may be solving $A^\star x = b^\star$, where the entries $A^\star$ and $b^\star$ have been subject to noise and inaccuracy. It is important to know how those translate to errors in determining x.

Vector and Matrix 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 $x \in \mathbb{R}^n$, we symbolize its norm by $||x||$. The defining properties of a norm are:

1. $||x|| \geq 0$ for all $x \in \mathbb{R}^n$, also $||x|| = 0$ if and only if $x = 0$.
2. $||\alpha x|| = |\alpha| \cdot ||x||$ for all $\alpha \in \mathbb{R}, x \in \mathbb{R}^n$.
3. $||x + y|| \leq ||x|| + ||y||$ for all $x, y \in \mathbb{R}^n$.

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.

Vector norms - why are they needed?

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

When solving a system of linear equations $Ax = b$, the exact solution x_{exact} as well as any approximation x_{approx} are vectors, and the error $e = x_{\mathsf{approx}} - x_{\mathsf{exact}}$ is a vector too! It is not as straightforward to assess the "magnitude" of such a vector-valued error. For e.g., consider $e_1, e_2 \in \mathbb{R}^{1000}$, and

$$e_1 = \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ \vdots \\ 0.1 \end{pmatrix}, e_2 = \begin{pmatrix} 100 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

Which one is worse? e_1 has a modest amount of error, distributed over all components. In 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. Some norms which satisfy the properties of vector norms are (here, $x = (x_1, x_2, \ldots, x_n)$):

1. The L_1-norm (or 1-norm)

$$||x||_1 = \sum_{i=1}^{n} |x_i|$$

2. The L_2 norm (or 2-norm, or Euclidean norm)

$$||x||_2 = \sqrt{\sum_{i=1}^{n} x_i^2}$$

3. The infinity norm (or max-norm)

$$||x||_\infty = \max_{1 \le i \le n} |x_i|$$

4. (Less common) L_p norm

$$||x||_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}$$

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

- $||x||_1 = \sum_{i=1}^n |x_i| \geq 0$
- if $x = 0 \Rightarrow \sum_{i=1}^n |x_i| = 0 \Rightarrow |x_i| = 0, \forall i \Rightarrow x = 0$
- $||\alpha x|| = \sum_{i=1}^n |\alpha x_i| = |\alpha| \sum_{i=1}^n |x_i| = |\alpha| \cdot ||x||_1$
- $||x + y|| = \sum_{i=1}^n |x_i + y_i| \leq \sum_{i=1}^n |x_i| + \sum_{i=1}^n |y_i| = ||x||_1 + ||y||_1$

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

We can actually define norms for (square) matrices as well. A matrix norm is a function $|| \cdot || : \mathbb{R}^{n \times n} \to \mathbb{R}$ which satisfies:

1. $||M|| \geq 0, \forall M \in \mathbb{R}^{n \times n}$, $||M|| = 0$ if and only if $M = O$.
2. $||\alpha M|| = |\alpha| \cdot ||M||$
3. $||M + N|| \leq ||M|| + ||N||$
4. $||M \cdot N|| \leq ||M|| \cdot ||N||$

(Property (4) is the one that has slightly different flavor than vector norms.)

Although *more types of matrix norms exist*, one common category is that of matrix norms *induced* by vector norms.

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

$$||M||_\star = \max_{x \in \mathbb{R}^n, x \neq 0} \frac{||Mx||_\star}{||x||_\star}$$

or equivalently,

$$||M||_\star = \max_{x \in \mathbb{R}^n, ||x||_\star = 1} ||Mx||_\star$$

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

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

We can easily show that induced norms satisfy properties (1) through (4). Properties (1)-(3) are rather trivial, e.g.:

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

Property (4) is slightly trickier to show. First, a lemma:

Lemma 1. *If* $||\cdot||$ *is a matrix norm induced by a vector norm* $||\cdot||$*, then:*

$$||Ax|| \leq ||A|| \cdot ||x||$$

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

$$||A|| = \max_{x\neq 0} \frac{||Ax||}{||x||} \geq \frac{||Ay||}{||y||} \Rightarrow ||Ay|| \leq ||A|| \cdot ||y||$$

This holds for $y \neq 0$, but we can see it is also true for $y = 0$. □

Property (4)

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

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):

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

($||\cdot||_2$ is much more complicated!) Where do these vector/matrix norms come handy?