CS 450: Numerical Anlaysis

Lecture 5

Chapter 2 – Linear Systems

Solving Linear Systems

Edgar Solomonik

Department of Computer Science
University of Illinois at Urbana-Champaign

February 2, 2018

Solving Triangular Systems

- **$Lx = b$ if L is lower-triangular is solved by forward substitution:**

$$\begin{aligned} l_{11}x_1 &= b_1 \\ l_{21}x_1 + l_{22}x_2 &= b_2 \\ l_{31}x_1 + l_{32}x_2 + l_{33}x_3 &= b_2 \\ &\vdots \end{aligned} \quad \Rightarrow \quad \begin{aligned} x_1 &= b_1/l_{11} \\ x_2 &= (b_2 - l_{21}x_1)/l_{22} \\ x_3 &= (b_3 - l_{31}x_1 - l_{32}x_2)/l_{33} \\ &\vdots \end{aligned}$$

- **Computational complexity of forward/backward substitution:**
  The total cost for $\boldsymbol{x} \in \mathbb{R}^n$ is $n^2/2$ multiplications and $n^2/2$ additions to leading order. So the asymptotic complexity is $O(n^2)$, the same as for a matrix-vector product.

Solving Triangular Systems

- **Existence of solution to $Lx = b$:**
  If some $l_{ii} = 0$, the solution may not exist, and $\boldsymbol{L}^{-1}$ does not exist.
- **Invertibility of L and existence of solution:**
  Even if some $l_{ii} = 0$ and $\boldsymbol{L}^{-1}$ does not exist, the system may have a solution. The solution will not be unique since columns of $\boldsymbol{L}$ are necessarily linearly dependent if a diagonal element is zero.

Properties of Triangular Matrices

- **$Z = XY$ is lower triangular is X and Y are both lower triangular:**

  Holds trivially when $n = 1$, then for $n > 1$,

$$\begin{bmatrix} \boldsymbol{Z}_{11} & \boldsymbol{Z}_{12} \\ \boldsymbol{Z}_{21} & \boldsymbol{Z}_{22} \end{bmatrix} = \begin{bmatrix} \boldsymbol{X}_{11} & \\ \boldsymbol{X}_{21} & \boldsymbol{X}_{22} \end{bmatrix} \begin{bmatrix} \boldsymbol{Y}_{11} & \\ \boldsymbol{Y}_{21} & \boldsymbol{Y}_{22} \end{bmatrix}.$$

  By induction $\boldsymbol{Z}_{11} = \boldsymbol{X}_{11}\boldsymbol{Y}_{11}$ and $\boldsymbol{Z}_{22} = \boldsymbol{X}_{22}\boldsymbol{Y}_{22}$ are lower-triangular. Then it suffices to observe that $\boldsymbol{Z}_{12} = \boldsymbol{0}$.

- **$\boldsymbol{L}^{-1}$ is lower triangular if it exists:**

  We give a constructive proof by providing an algorithm for triangular matrix inversion, We need $\boldsymbol{Y} = \boldsymbol{X}^{-1}$ so

$$\begin{bmatrix} \boldsymbol{Y}_{11} & \\ \boldsymbol{Y}_{21} & \boldsymbol{Y}_{22} \end{bmatrix} \begin{bmatrix} \boldsymbol{X}_{11} & \\ \boldsymbol{X}_{21} & \boldsymbol{X}_{22} \end{bmatrix} = \begin{bmatrix} \boldsymbol{I} & \\ & \boldsymbol{I} \end{bmatrix},$$

  from which we can deduce

$$\boldsymbol{Y}_{11} = \boldsymbol{X}_{11}^{-1}, \quad \boldsymbol{Y}_{22} = \boldsymbol{X}_{22}^{-1}, \quad \boldsymbol{Y}_{21} = -\boldsymbol{Y}_{22}\boldsymbol{X}_{21}\boldsymbol{Y}_{11}.$$

LU Factorization

- **An *LU factorization* consists of a unit-diagonal lower-triangular *factor* $\boldsymbol{L}$ and upper-triangular factor $\boldsymbol{U}$ such that $\boldsymbol{A} = \boldsymbol{L}\boldsymbol{U}$:**
  - *Unit-diagonal implies each $l_{ii} = 1$, leaving $n(n-1)/2$ unknowns in $\boldsymbol{L}$ and $n(n+1)/2$ unknowns in $\boldsymbol{U}$, for a total of n^2, the same as the size of $\boldsymbol{A}$.*
  - *Once we have an LU factorization of $\boldsymbol{A}$, we can solve the linear system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$*
    1. *using forward substitution $\boldsymbol{L}\boldsymbol{y} = \boldsymbol{b}$*
    2. *using backward substitution to solve $\boldsymbol{U}\boldsymbol{x} = \boldsymbol{y}$*
  - *Assuming invertibility, this corresponds to computing $\boldsymbol{x} = \boldsymbol{U}^{-1}\boldsymbol{L}^{-1}\boldsymbol{b}$*
  - *It suffices to have a forward substitution routine and reversing the ordering of vector elements, done by taking the product with $\boldsymbol{P}$, $p_{ij} = \delta(i, n+1-j)$,*

$$\boldsymbol{L}\boldsymbol{U} = \boldsymbol{L}\boldsymbol{P}\underbrace{\boldsymbol{P}\boldsymbol{U}\boldsymbol{P}}_{\tilde{\boldsymbol{L}}}\boldsymbol{P} = \boldsymbol{L}\boldsymbol{P}\tilde{\boldsymbol{L}}\boldsymbol{P}$$

  - *For rectangular matrices $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, one can consider a full LU factorization, with $\boldsymbol{L} \in \mathbb{R}^{m \times \max(m,n)}$ and $\boldsymbol{U} \in \mathbb{R}^{\max(m,n) \times n}$, but it is fully described by a reduced LU factorization, with $\boldsymbol{L} \in \mathbb{R}^{m \times \min(m,n)}$ and $\boldsymbol{U} \in \mathbb{R}^{\min(m,n) \times n}$.*

Gaussian Elimination

- **The LU factorization may not exist:** Consider matrix $\begin{bmatrix} 3 & 2 \\ 6 & 4 \\ 0 & 3 \end{bmatrix}$.

  We can infer what the first row of $\boldsymbol{L}$ and column of $\boldsymbol{U}$ directly from the matrix,

  $$\begin{bmatrix} 3 & 2 \\ 6 & 4 \\ 0 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 2 & 1 \\ 0 & l_{32} \end{bmatrix} \begin{bmatrix} 3 & 2 \\ 0 & u_{21} \end{bmatrix}.$$

  Then we can observe that $4 = 4 + u_{21}$ so $u_{21} = 0$, but at the same time $l_{32}u_{21} = 3$, which is a contradiction.
  More generally, if for any partitioning $\begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix}$ the leading minor is singular ($\det(\boldsymbol{A}_{11}) = 0$), $\boldsymbol{A}$ has no LU factorization.

- **Permutation of variables enables us to transform the linear system so the LU factorization does exist:**

  If $\boldsymbol{A}$ is not singular, its leading k columns for any k have a span of dimension k, and so a permutation of their rows exists so $\boldsymbol{A}_{11} \in \mathbb{R}^{k \times k}$ is not singular.

Gaussian Elimination Algorithm

- **Algorithm for factorization is derived from equations given by $A = LU$:**

$$\begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix} = \begin{bmatrix} \boldsymbol{L}_{11} & \\ \boldsymbol{L}_{21} & \boldsymbol{L}_{22} \end{bmatrix} \begin{bmatrix} \boldsymbol{U}_{11} & \boldsymbol{U}_{12} \\ & \boldsymbol{U}_{22} \end{bmatrix}$$

  - *Obtain LU factorization of the leading minor $\boldsymbol{A}_{11} = \boldsymbol{L}_{11}\boldsymbol{U}_{11}$ by recursion*
  - *Solve sets of triangular linear systems can be solved to obtain $\boldsymbol{L}_{21}$ and $\underset{\sim}{U}_{12}$ from $\boldsymbol{A}_{21} = \boldsymbol{L}_{21}\boldsymbol{U}_{11}$ and $\boldsymbol{A}_{12} = \boldsymbol{L}_{11}\boldsymbol{U}_{12}$*
  - *Obtain $\boldsymbol{L}_{22}$ and $\boldsymbol{U}_{22}$ by recursion of LU on Schur complement*

$$\boldsymbol{S} = \boldsymbol{A}_{22} - \boldsymbol{L}_{21}\boldsymbol{U}_{12} = \boldsymbol{A}_{22} - \boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12} = \boldsymbol{L}_{22}\boldsymbol{U}_{22}$$

- **The kth column of L is given by the kth *elementary matrix* M_k:**

$$\boldsymbol{M}_k \begin{bmatrix} v_1 & \cdots & v_k & 0 & \cdots & 0 \end{bmatrix}^T = \boldsymbol{v}$$

Elimination Matrices

- **An elimination matrix M_k satisfies the following properties:**
  - *It is a rank-1 perturbation of the identity that is unit-diagonal and lower-triangular,*

$$\boldsymbol{M}_k = \boldsymbol{I} - \boldsymbol{m}_k \boldsymbol{e}_k^T = \boldsymbol{I} - \begin{bmatrix} \tilde{\boldsymbol{m}}_k \\ \boldsymbol{0} \end{bmatrix} \boldsymbol{e}_k^T$$

  - *It reduces a given vector to its first k elements*

$$\boldsymbol{M}_k \begin{bmatrix} a_1 \\ \vdots \\ a_k \\ 0 \\ \vdots \end{bmatrix} = \boldsymbol{a}$$

  - $\boldsymbol{M}_k^{-1} = \boldsymbol{I} + \boldsymbol{m}_k \boldsymbol{e}_k^T = 2\boldsymbol{I} - \boldsymbol{M}_k$
  - $\boldsymbol{M}_j \boldsymbol{M}_k = (\boldsymbol{I} - \boldsymbol{m}_j \boldsymbol{e}_j^T)(\boldsymbol{I} - \boldsymbol{m}_k \boldsymbol{e}_k^T) = (\boldsymbol{I} - \begin{bmatrix} \boldsymbol{m}_j & \boldsymbol{m}_k \end{bmatrix} \begin{bmatrix} \boldsymbol{e}_j & \boldsymbol{e}_k \end{bmatrix}^T) = \boldsymbol{M}_j + \boldsymbol{M}_k - \boldsymbol{I}$

Gaussian Elimination with Partial Pivoting

- ***Partial pivoting*** **permutes rows to make divisor u_{ii} is maximal at each step:**

  Based on our argument above, for any matrix $\boldsymbol{A}$ there exists a permutation matrix $\boldsymbol{P}$ that can permute the rows of $\boldsymbol{A}$ to permit an LU factorization,

  $$\boldsymbol{PA} = \boldsymbol{LU}.$$

  Partial pivoting finds such a permutation matrix $\boldsymbol{P}$ one row at a time. The ith row is selected to maximize the magnitude of the leading element (over elements in the first column), which becomes the entry u_{ii}. This selection ensures that we are never forced to divide by zero during Gaussian elimination and that the magnitude of any element in $\boldsymbol{L}$ is at most 1.

- **A row permutation corresponds to an application of a *row permutation matrix* $\boldsymbol{P}_{jk} = \boldsymbol{I} - (\boldsymbol{e}_j - \boldsymbol{e}_k)(\boldsymbol{e}_j - \boldsymbol{e}_k)^T$:**

  If we permute row i_j to be the leading (ith) row at the ith step, the overall permutation matrix is given by $\boldsymbol{P}^T = \prod_{i=1}^{n-1} \boldsymbol{P}_{ii_j}$, generally, $p_{ij} \neq \delta(j, i_j)$.

Complete Pivoting and Error Bounds

- ***Complete pivoting*** **permutes rows and columns to make divisor u_{ii} is maximal at each step:**

  Partial pivoting bounds the size of elements in $\boldsymbol{L}$, but elements in $\boldsymbol{U}$ can have magnitudes larger than elements of $\boldsymbol{A}$. Complete pivoting bounds this growth, by ensuring that

- **For LU, the backward error δA, so that $\hat{\boldsymbol{L}}\hat{\boldsymbol{U}} = \boldsymbol{A} + \delta \boldsymbol{A}$, satisfies bound $|\delta a_{ij}| \leq \epsilon(|\hat{\boldsymbol{L}}| \cdot |\hat{\boldsymbol{U}}|)_{ij}$:**

  For an arbitrary a_{ij}, consider the partitioning $\begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix}$ where $\boldsymbol{A}_{11}$ is of dimension $\min(i,j)-1$. After the Schur complement update $\boldsymbol{S} = fl(\boldsymbol{A}_{22} - \hat{\boldsymbol{L}}_{21}\hat{\boldsymbol{U}}_{12})$, the entry in $\boldsymbol{S}$ corresponding to a_{ij} (an entry in $\boldsymbol{A}_{22}$) will become an entry of $\boldsymbol{U}$ or an entry of $\hat{\boldsymbol{L}}$ (after a division that can only shrink the error). Thus the $\hat{\mathbf{L}}$ and $\hat{\mathbf{U}}$ are factors of a matrix $\boldsymbol{A} + \delta\boldsymbol{A}$ where the perturbation is bounded by the error of the inner product necessary to compute any Schur complement entry, so $|\delta a_{ij}| \leq \epsilon(|\hat{\boldsymbol{L}}| \cdot |\hat{\boldsymbol{U}}|)_{ij}$.