

Determining an LU factorization



Demo: LU Factorization [cleared]

Computational Cost

What is the computational cost of multiplying two $n \times n$ matrices?

$$O(n^3)$$

▶ $u_{11} = a_{11}, \mathbf{u}_{12}^T = \mathbf{a}_{12}^T.$

▶ $l_{21} = \mathbf{a}_{21}/u_{11}.$

▶ $L_{22}U_{22} = A_{22} - l_{21}\mathbf{u}_{12}^T.$

0 "flops"
 ~~$O(n)$~~ "stray"
 $O(n^2)$



What is the computational cost of carrying out LU factorization on an $n \times n$ matrix?

$$O(n^3)$$

Demo: Complexity of Mat-Mat multiplication and LU [cleared]

LU: Failure Cases?

Is LU/Gaussian Elimination bulletproof?

$$A = \begin{pmatrix} 0 & 1 \\ 2 & 1 \end{pmatrix}$$

$$\begin{pmatrix} 1 & & \\ l_{21} & 1 & \end{pmatrix} \begin{pmatrix} u_{11} & u_{12} \\ 0 & 1 \\ 2 & 1 \end{pmatrix}$$

$$u_{11} = 0$$

$$\underbrace{l_{21} \cdot u_{11}}_0 + 0 \cdot 1 = 2$$

Saving the LU Factorization

What can be done to get something *like* an LU factorization?

Idea: control the magnitude
of the denominator is



L , make as big as possible.

swap rows \rightarrow partial pivoting

swap rows and cols \rightarrow complete pivoting

Demo: LU Factorization with Partial Pivoting [cleared]

Cholesky: LU for Symmetric Positive Definite

LU can be used for SPD matrices. But can we do better?

$$A = LL^T$$
$$\begin{pmatrix} d_{11} \\ \vec{l}_{21} \\ L_{22} \end{pmatrix} \begin{pmatrix} d_{11} & \vec{l}_{21}^T \\ & L_{22} \end{pmatrix} = \begin{pmatrix} A \\ \dots \end{pmatrix}$$
$$d_{11}^2 = a_{11} \geq 0 \Leftrightarrow A \text{ SPSPD}$$

↳ semi

- cheaper than LU
- no need for pivoting

More cost concerns

$$\underbrace{LU}_{y} x = b$$

What's the cost of solving $Ax = b$?

$$O(n^3) \text{ for LU, } O(n^2) \text{ for fw/bw}$$

What's the cost of solving $Ax = b_1, b_2, \dots, b_n$?

$$O(n^3)$$

What's the cost of finding A^{-1} ?

$$Ax = b \quad \underbrace{Ax = I}_{x = A^{-1}b} \quad \Leftrightarrow \quad \begin{matrix} Ax_1 = e_1 \\ \vdots \\ Ax_n = e_n \end{matrix}$$

Cost: Worrying about the Constant, BLAS

$O(n^3)$ really means

$$\alpha \cdot n^3 + \beta \cdot n^2 + \gamma \cdot n + \delta.$$

All the non-leading and constants terms swept under the rug. But: at least the leading constant ultimately matters.

Shrinking the constant: surprisingly hard (even for 'just' matmul)

Idea: Rely on library implementation: *BLAS* (Fortran)

Level 1 $\mathbf{z} = \alpha \mathbf{x} + \mathbf{y}$ vector-vector operations

$O(n)$

?axpy

Level 2 $\mathbf{z} = \mathbf{A}\mathbf{x} + \mathbf{y}$ matrix-vector operations

$O(n^2)$

?gemv

Level 3 $\mathbf{C} = \mathbf{A}\mathbf{B} + \beta \mathbf{C}$ matrix-matrix operations

$O(n^3)$

?gemm, ?trsm

$$[\mathbf{A}\mathbf{B}]_{ij} = \sum_k A_{ik} B_{kj}$$

Show (using perf): numpy matmul calls BLAS dgemm

LAPACK

LAPACK: Implements 'higher-end' things (such as LU) using BLAS
Special matrix formats can also help save const significantly, e.g.

- ▶ banded ba
- ▶ sparse sp
- ▶ symmetric sy
- ▶ triangular tr
- ▶ - *general* ge

Sample routine names:

- ▶ dgesvd, zgesdd?
- ▶ dgetrf, dgetrs

LU on Blocks: The Schur Complement

Given a matrix

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}, = \begin{pmatrix} B & D \end{pmatrix} \begin{pmatrix} A & \\ & D \end{pmatrix}$$

can we do 'block LU' to get a *block triangular matrix*?

Multiply top row $-CA^{-1}$, add to second

$$\begin{pmatrix} A & B \\ 0 & D - CA^{-1}B \end{pmatrix}$$

"Schur complement"

LU: Special cases

$$P \mathbf{0} = \underbrace{\mathbf{I}}_L \mathbf{0}$$

What happens if we feed a non-invertible matrix to LU?

$$\underbrace{P}_\checkmark \underbrace{A}_\times = \underbrace{L}_\checkmark \underbrace{U}_\times$$

What happens if we feed LU an $m \times n$ non-square matrices?

Round-off Error in LU without Pivoting

Consider factorization of $\begin{bmatrix} \epsilon & 1 \\ 1 & 1 \end{bmatrix}$ where $\epsilon < \epsilon_{\text{mach}}$:

.