

Cost of LU factorization

```
## Algorithm 1
## Factorization using the block-format,
## creating new matrices L and U
## and not modifying A
print("LU factorization using Algorithm 1")
L = np.zeros((n,n))
U = np.zeros((n,n))
M = A.copy()
for i in range(n):
    U[i,i:] = M[i,i:]
    L[i:,i] = M[i:,i]/U[i,i]
    M[i+1:,i+1:] -= np.outer(L[i+1:,i],U[i,i+1:])
```

Side note:

$$\sum_{i=1}^m i = \frac{1}{2}m(m+1)$$

$$\sum_{i=1}^m i^2 = \frac{1}{6}m(m+1)(2m+1)$$

Number of divisions: $(n-1) + (n-2) + \dots + 1 = n(n-1)/2$


Number of multiplications $(n-1)^2 + (n-2)^2 + \dots + (1)^2 = \frac{n^3}{3} - \frac{n^2}{2} + \frac{n}{6}$

Number of subtractions: $(n-1)^2 + (n-2)^2 + \dots + (1)^2 = \frac{n^3}{3} - \frac{n^2}{2} + \frac{n}{6}$



Computational complexity is $O(n^3)$

Cost of LU factorization



$(n-1)$ divisions

$M = M$ - outer (l_{21}, u_{21})

$(n-1)^2$ multiplications

$(n-1)^2$ subtractions

TOTAL :

$$\frac{2n^3}{3} - n^2 + \frac{n}{3} + \frac{n(n-1)}{2}$$

Complexity $n \rightarrow \infty$

$$O(n^3)$$

$$\text{divisions} = (n-1) + (n-2) + (n-3) + \dots + 1 = \frac{n(n-1)}{2}$$

$$\begin{aligned} \text{multip. or subtractions} &= (n-1)^2 + (n-2)^2 + \dots + 1 = \sum_{i=1}^{n-1} i^2 = \frac{n^3}{3} - \frac{n^2}{2} + \frac{n}{6} \end{aligned}$$

LU factorization: $\frac{2}{3}n^3$; Cholesky: $\frac{1}{3}n^3$

Clicker question

Which of the following statements are true about the LU factorization of an $n \times n$ matrix A , assuming LU factorization of A exists and not considering any row/column interchanges?

Select all that apply:

- 1) $A = LU$. **True**
- 2) LU factorization is exactly performing Gaussian elimination. **True**
- 3) We can solve for $LUx = b$ instead of solving $Ax = b$ to obtain x . **True**
- 4) L is a lower triangular matrix, and is exactly the lower part of A but with unit diagonal. **False**
- 5) U is an upper triangular matrix, and is exactly the upper part of A (including diagonal). **False**

A) 1, 2, 3

B) 1, 2, 3, 5

C) 1, 3

D) 1, 2, 3, 4, 5

E) 4, 5

Solving linear systems

In general, we can solve a linear system of equations following the steps:

- 1) Factorize the matrix \mathbf{A} : $\mathbf{A} = \mathbf{LU}$ (complexity $O(n^3)$)
- 2) Solve $\mathbf{L}\mathbf{y} = \mathbf{b}$ (complexity $O(n^2)$)
- 3) Solve $\mathbf{U}\mathbf{x} = \mathbf{y}$ (complexity $O(n^2)$)

But why should we decouple the factorization from the actual solve?

(Remember from Linear Algebra, Gaussian Elimination does not decouple these two steps...)

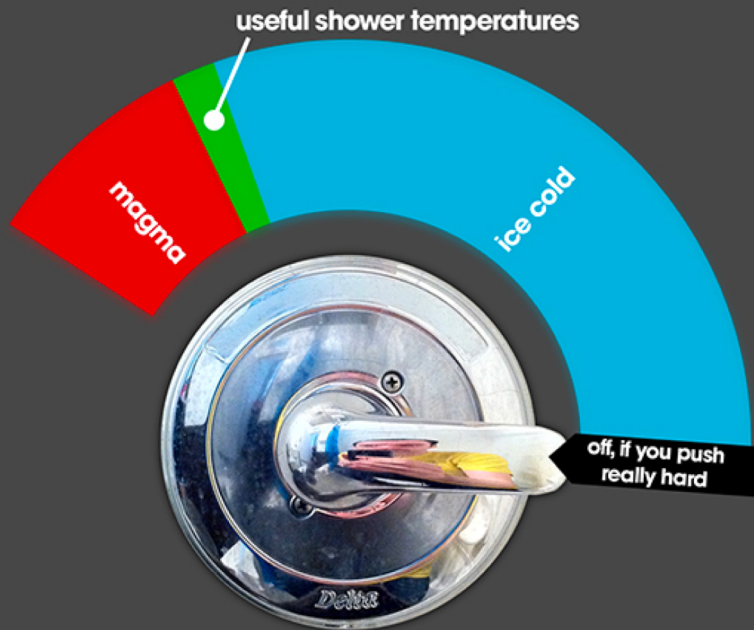
Think about solving

$\underline{\mathbf{A}}\underline{\mathbf{x}} = \underline{\mathbf{b}}$ for different
right-hand sides $\underline{\mathbf{b}}$

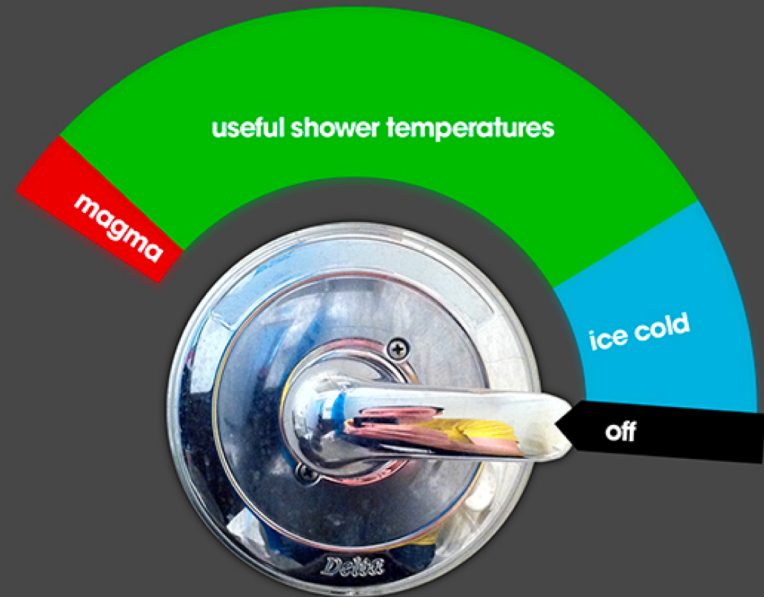
Linear System of Equations - Conditioning

the shower faucet

how they are:



how they should be:



WHAT IT LOOKS LIKE



WHAT IT FEELS LIKE



Numerical experiments

Input has uncertainties:

- Errors due to representation with finite precision
- Error in the sampling


Once you select your numerical method , how much error should you expect to see in your **output**?

Is your method sensitive to errors (perturbation) in the input?

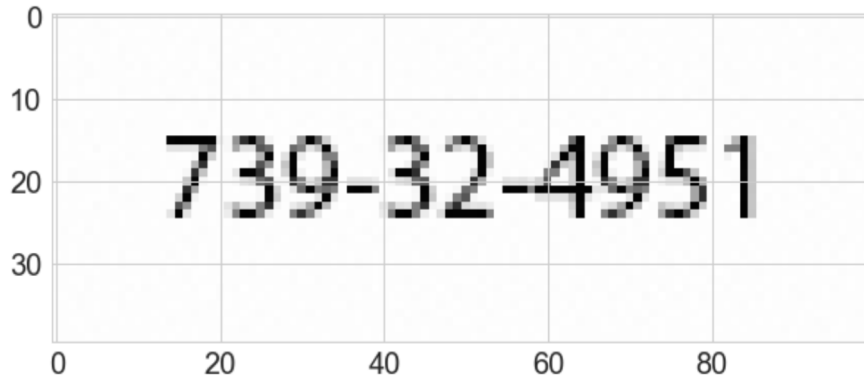
Demo “HilbertMatrix-ConditionNumber”

Demo “image-blur-inverse-v3”

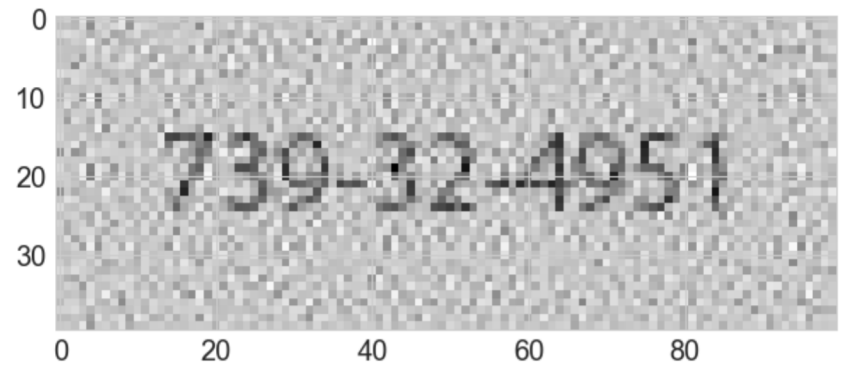
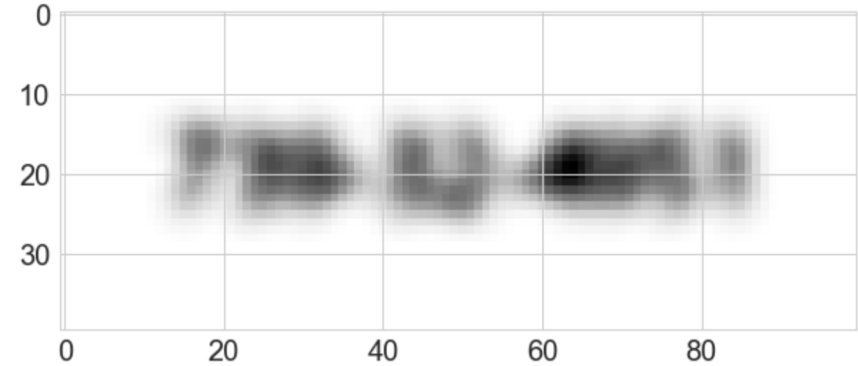
$$\mathbf{b} + a * 10^{-6} \quad (a \in (0,1))$$



Solve $\mathbf{A} \mathbf{x} = \mathbf{b}$ for \mathbf{x}



$$\mathbf{b} + a * 10^{-4} \quad (a \in (0,1))$$



Is your method sensitive to errors (perturbation) in the input?

How much noise can we add to the input data?


How can we define "little" amount of noise? Should be relative with the magnitude of the data.

Sensitivity of Solutions of Linear Systems

Suppose we start with a non-singular system of linear equations $A \mathbf{x} = \mathbf{b}$.

We change the right-hand side vector \mathbf{b} (input) by a small amount $\Delta \mathbf{b}$.

How much the solution \mathbf{x} (output) changes, i.e., how large is $\Delta \mathbf{x}$?



$$\hat{x} = x + \Delta x \quad ; \quad \hat{b} = b + \Delta b \quad ; \quad \hat{A} = A + \Delta A$$

$$\frac{\text{output er}}{\text{input er}} = \frac{\|\Delta x\| / \|x\|}{\|\Delta b\| / \|b\|} = \frac{\|\Delta x\| \|b\|}{\|\Delta b\| \|x\|}$$

$$\frac{\text{output er}}{\text{input er}} \leq \underbrace{\|A^{-1}\| \|A\|}_{\text{cond}(A)}$$

$$Ax = b \Rightarrow A(x + \Delta x) = (b + \Delta b) \Rightarrow \cancel{Ax} + A\Delta x = \cancel{b} + \Delta b$$

$$A \Delta x = \Delta b$$

$$= \frac{\|A^{-1} \Delta b\| \|b\|}{\|\Delta b\| \|x\|} \leq \frac{\|A^{-1}\| \|\Delta b\| \|b\|}{\cancel{\|b\|} \|x\|} = \frac{\|A^{-1}\| \|Ax\|}{\|x\|} \leq \frac{\|A^{-1}\| \|A\| \|x\|}{\cancel{\|x\|}}$$

Condition number

The condition number is a measure of sensitivity of solving a linear system of equations to variations in the input.

$$\text{output } e_r \leq \text{cond}(A) \text{ input } e_r$$

$$\text{cond}(A) = \|A^{-1}\| \|A\|$$

$$\frac{\|\Delta x\|}{\|x\|} \leq \text{cond}(A) \frac{\|\Delta b\|}{\|b\|}$$

we should be specific and

say $\text{cond}_p(A) = \|A^{-1}\|_p \|A\|_p$

here, when we omit, we use $\text{cond}_2(A)$.

systems with large $\text{cond}(A)$ can have large errors in the output even from small changes in the input (shower example)

Clicker question

Give an example of a matrix that is very well-conditioned (i.e., has a condition number that is good for computation). Select the best possible condition number(s) of a matrix?

- ~~A) $\text{cond}(A) < 0$~~
~~B) $\text{cond}(A) = 0$~~
~~C) $0 < \text{cond}(A) < 1$~~
D) $\text{cond}(A) = 1$
~~E) $\text{cond}(A) = \text{large numbers}$~~
- $\text{cond}(A) = \|A^{-1}\| \|A\|$ and $\|A\| > 0$
 $\rightarrow \text{cond}(A) = \|A^{-1}\| \|A\| \geq \|A^{-1}A\|$ (triangle inequality)
 $= \|I\| = 1$
 $\Rightarrow \text{cond}(A) \geq 1$!
 \rightarrow ill-conditioned matrices!

Condition number

$$\frac{\|\Delta \mathbf{x}\|}{\|\mathbf{x}\|} \leq \text{cond}(\mathbf{A}) \frac{\|\Delta \mathbf{b}\|}{\|\mathbf{b}\|}$$

Small condition numbers mean not a lot of error amplification. Small condition numbers are good!

The identity matrix should be well-conditioned:

$$\|\mathbf{I}\| = \max_{\|\mathbf{x}\|=1} \|\mathbf{I} \mathbf{x}\| = 1$$

It turns out that this is the smallest possible condition number:

$$\text{cond}(\mathbf{A}) = \|\mathbf{A}^{-1}\| \|\mathbf{A}\| \geq \|\mathbf{A}^{-1} \mathbf{A}\| = \|\mathbf{I}\| = 1$$

If \mathbf{A}^{-1} does not exist, then $\text{cond}(\mathbf{A}) = \infty$ (by convention)

Recall Induced Matrix Norms

$$\|\mathbf{A}\|_1 = \max_j \sum_{i=1}^n |A_{ij}|$$

Maximum absolute column sum of the matrix \mathbf{A}

$$\|\mathbf{A}\|_\infty = \max_i \sum_{j=1}^n |A_{ij}|$$

Maximum absolute row sum of the matrix \mathbf{A}

$$\|\mathbf{A}\|_2 = \max_k \sigma_k$$

σ_k are the singular value of the matrix \mathbf{A}

Clicker question

Condition Number of a Diagonal Matrix

What is the 2-norm-based condition number of the diagonal matrix

$$A = \begin{bmatrix} 100 & 0 & 0 \\ 0 & 13 & 0 \\ 0 & 0 & 0.5 \end{bmatrix} ?$$

A) 1

B) 50

C) 100

D) 200

$$\|A^{-1}\| = 2$$

$$\|A\| = 100$$

$$\text{cond}(A) = 200$$

About condition numbers

1. For any matrix \mathbf{A} , $\text{cond}(\mathbf{A}) \geq 1$
2. For the identity matrix \mathbf{I} , $\text{cond}(\mathbf{I}) = 1$
3. For any matrix \mathbf{A} and a nonzero scalar γ , $\text{cond}(\gamma\mathbf{A}) = \text{cond}(\mathbf{A})$
4. For any diagonal matrix \mathbf{D} , $\text{cond}(\mathbf{D}) = \frac{\max|d_i|}{\min|d_i|}$
5. The condition number is a measure of how close a matrix is to being singular: a matrix with large condition number is nearly singular, whereas a matrix with a condition number close to 1 is far from being singular
6. The determinant of a matrix is NOT a good indicator is a matrix is near singularity

Clicker question

The need for pivoting depends on whether the matrix is singular.

pivoting is needed even for nonsingular matrices

A) True

B) False

(* poorly written.

Intended to say "matrix is well conditioned if $\text{cond}(A) \leq 1$ "

which is a False statement, since $\text{cond}(A) \geq 1$.

Which of the following statements is correct?

Choice*

A) A singular matrix does not have a solution False (also infinitely many solutions)

B) A matrix is well conditioned if its condition number is less or equal to 1 (*)

C) A nonsingular matrix always has a solution True

D) 1-norm of a matrix is the absolute column sum False

largest value from

Condition Number of Orthogonal Matrices

What is the 2-norm condition number of an orthogonal matrix A ?

$$\mathit{cond}(A) = \|A^{-1}\|_2 \|A\|_2 = \|A^T\|_2 \|A\|_2 = 1$$

That means orthogonal matrices have optimal conditioning.

They are very well-behaved in computation.

Residual versus error

Our goal is to find the solution \mathbf{x} to the linear system of equations $\mathbf{A} \mathbf{x} = \mathbf{b}$

Let us recall the solution of the perturbed problem

$$\hat{\mathbf{x}} = \mathbf{x} + \Delta \mathbf{x}$$

which can be
solution of

$$\begin{aligned} \mathbf{A} \hat{\mathbf{x}} &= \mathbf{b} + \Delta \mathbf{b}; & (\mathbf{A} + \Delta \mathbf{A}) \hat{\mathbf{x}} &= \mathbf{b}; \\ & & (\mathbf{A} + \Delta \mathbf{A}) \hat{\mathbf{x}} &= \mathbf{b} + \Delta \mathbf{b} \end{aligned}$$

error $= \|\mathbf{x} - \hat{\mathbf{x}}\| \rightarrow$ but we don't know $\|\mathbf{x}\|$!

Let's define the residual vector

$$\underline{\mathbf{r}} = \underline{\mathbf{A}} \hat{\underline{\mathbf{x}}} - \underline{\mathbf{b}}$$

we want to minimize
the error, but we only
know how to measure
the residual ...

Residual versus error

By solving $Ax = b$ with LU factorization, the residual $\underline{r} = A\hat{x} - b$ satisfies the following:

$$\frac{\|r\|}{\|A\| \|\hat{x}\|} \leq c \epsilon_m$$

Where c is large without pivoting and small with partial pivoting.

Therefore, Gaussian elimination with partial pivoting yields **small relative residual regardless of conditioning of the system.**

We now know that the residual norm is bounded by small value (when using partial pivoting). But can we ensure the error $\Delta x = x - \hat{x}$ is also small?

Residual versus error

Let us first obtain the norm of the error:

$$\|x - \hat{x}\| = \|A^{-1}b - A^{-1}A\hat{x}\| = \|A^{-1}(b - A\hat{x})\| = \|A^{-1}r\|$$

$$\|\Delta x\| = \|A^{-1}r\|$$

$$\frac{\|\Delta x\|}{\|x\|} = \frac{\|A^{-1}r\|}{\|x\|} \leq \frac{\|A^{-1}\| \|r\|}{\|x\|} = \frac{\|A^{-1}\| \|A\| \|r\|}{\|A\| \|x\|}$$

$$\frac{\|\Delta x\|}{\|x\|} \leq \text{cond}(A) \frac{\|r\|}{\|A\| \|x\|}$$

if $\|r\|$ is small and
 A is well-conditioned

\Downarrow
 $\|e_r\|$ is small

Clicker question

When solving a system of linear equations via LU with partial pivoting, which of the following is guaranteed to be small?

A) Relative residual: $\frac{\|r\|}{\|A\|\|x\|}$

B) Relative error: $\frac{\|\Delta x\|}{\|x\|}$

→ when is this true?
when $\text{cond}(A)$ is
also small!

C) Neither one of them

D) Both of them

Rule of thumb for conditioning

Suppose we want to find the solution \mathbf{x} to the linear system of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ using LU factorization with partial pivoting and backward/forward substitutions.

Suppose we compute the solution $\hat{\mathbf{x}}$.

If the entries in \mathbf{A} and \mathbf{b} are accurate to S decimal digits,

and $\text{cond}(\mathbf{A}) = 10^W$,

then the elements of the solution vector $\hat{\mathbf{x}}$ will be accurate to about

$$S - W$$

decimal digits

Rule of thumb

solving $A\hat{x} = b$

suppose A, b have X decimal digits of accuracy

and $\text{cond}(A) = 10^W$

then \hat{x} will have $Y = X - W$ decimal digits
of accuracy

Clicker question

Matrix Conditioning: Accurate digits

1 point

Let's say we want to solve the following linear system:

$$Ax = b$$

Assuming you are working with IEEE double precision floating point numbers, how many digits of accuracy will your answer have if $\kappa(A) = 1000$?

- A) 3
- B) 10
- C) 13**
- D) 16
- E) 32

double precision \rightarrow A, b have
16 digits of accuracy

the output (answer \hat{x}) will have
(16 - w) digits of accuracy

$$\text{cond}(A) = 10^3 \rightarrow w = 3 \rightarrow 16 - w = 13$$

Sparse Systems

Sparse Matrices

Some type of matrices contain many zeros.
Storing all those zero entries is wasteful!



How can we efficiently store large
matrices without storing tons of zeros?





- **Sparse matrices** (vague definition): matrix with few non-zero entries.
- For practical purposes: an $m \times n$ matrix is sparse if it has $O(\min(m, n))$ non-zero entries.
- This means roughly a constant number of non-zero entries per row and column.
- Another definition: “matrices that allow special techniques to take advantage of the large number of zero elements” (J. Wilkinson)

Sparse Matrices: Goals

- Perform standard matrix computations economically, i.e., without storing the zeros of the matrix.
- For typical Finite Element and Finite Difference matrices, the number of non-zero entries is $O(n)$



Sparse Matrices: MP example



Sparse Matrices

EXAMPLE:

Number of operations required to add two square dense matrices:

$$O(n^2)$$

Number of operations required to add two sparse matrices **A** and **B**:

$$O(\text{nnz}(\mathbf{A}) + \text{nnz}(\mathbf{B}))$$

where $\text{nnz}(\mathbf{X})$ = number of non-zero elements of a matrix **X**

Popular Storage Structures

DNS	Dense	ELL	Ellpack-Itpack
BND	Linpack Banded	DIA	Diagonal
COO	Coordinate	BSR	Block Sparse Row
CSR	Compressed Sparse Row	SSK	Symmetric Skyline
CSC	Compressed Sparse Column	BSR	Nonsymmetric Skyline
MSR	Modified CSR	JAD	Jagged Diagonal
LIL	Linked List		

note: CSR = CRS, CCS = CSC, SSK = SKS in some references

We will focus on COO and CSR!

Dense (DNS)

$$A = \begin{bmatrix} 0. & 1.9 & 0. & -5.2 \\ 0.3 & 0. & 9.1 & 0. \\ 4.4 & 5.8 & 3.6 & 0. \\ 0. & 0. & 7.2 & 2.7 \end{bmatrix}$$

$A_{shape} = (nrow, ncol)$

$$A_{dense} = [0. \quad 1.9 \quad 0. \quad -5.2 \quad 0.3 \quad 0. \quad 9.1 \quad 0. \quad 4.4 \quad 5.8 \quad 3.6 \quad 0. \quad 0. \quad 0. \quad 7.2 \quad 2.7]$$

Row 0 Row 1 Row 2 Row 3

- Simple
- Row-wise
- Easy blocked formats
- Stores all the zeros

Coordinate (COO)

total entries = 3 nnz

data \rightarrow float
col, row \rightarrow int

$$A = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 0. & 1.9 & 0. & -5.2 \\ 0.3 & 0. & 9.1 & 0. \\ 4.4 & 5.8 & 3.6 & 0. \\ 0. & 0. & 7.2 & 2.7 \end{bmatrix} \begin{matrix} 0 \\ 1 \\ 2 \\ 3 \end{matrix}$$

= $\underbrace{2 \text{ nnz}}_{\text{int}} + \underbrace{1 \text{ nnz}}_{\text{float}}$

$$\begin{aligned} \text{data} &= [1.9 \quad -5.2 \quad 0.3 \quad 9.1 \quad 4.4 \quad 5.8 \quad 3.6 \quad 7.2 \quad 2.7] \\ \text{col} &= [1 \quad 3 \quad 0 \quad 2 \quad 0 \quad 1 \quad 2 \quad 2 \quad 3] \\ \text{row} &= [0 \quad 0 \quad 1 \quad 1 \quad 2 \quad 2 \quad 2 \quad 3 \quad 3] \end{aligned}$$

- Simple
- Does not store the zero elements
- Not sorted
- *row* and *col*: array of integers
- *data*: array of doubles

can happen in any order:

$$\begin{aligned} \text{data} &= [9.1 \quad 5.8 \quad 7.2 \quad \dots] \\ \text{col} &= [2 \quad 1 \quad 2 \quad \dots] \\ \text{row} &= [1 \quad 2 \quad 3 \quad \dots] \end{aligned}$$

Example

Representing a Sparse Matrix in Coordinate (COO) Form

1 point

Consider the following matrix:

$$A = \begin{bmatrix} 0 & 0 & 1.3 \\ -1.5 & 0.2 & 0 \\ 5 & 0 & 0 \\ 0 & 0.3 & 3 \\ 0 & 0 & 0 \end{bmatrix}$$

- A) 56 bytes
- B) 72 bytes
- C) 96 bytes**
- D) 120 bytes
- E) 144 bytes

Suppose we store one row index (a 32-bit integer), one column index (a 32-bit integer), and one data value (a 64-bit float) for each non-zero entry in A . How many bytes in total are stored? Please note that 1 byte is equal to 8 bits.

store 6 floats + 12 integers

$$(6 \times 64 + 12 \times 32) \text{ bits} = 96 \text{ bytes}$$