

FIVE FACTORIZATIONS OF A MATRIX

Gilbert Strang (2023)

Introduction to Linear Algebra (6th ed)

KEY IDEAS AT THE START

Linear independence and dependence

$A\mathbf{x}$ = combination of columns of $A = x_1\mathbf{a}_1 + \cdots + x_n\mathbf{a}_n$

Each column of CR is a combination of the columns of C

$A = CR$: Column space from C , row space from R

1 $\mathbf{A} = \mathbf{CR}$ ($m \times r$) ($r \times n$) for any m by n matrix \mathbf{A}
of rank r

2 $\mathbf{A} = \mathbf{LU}$ or $\mathbf{PA} = \mathbf{LU}$ $L =$ lower triangular
 $U =$ upper triangular $P =$ permutation

For n by n invertible matrices, solve $\mathbf{Ax} = \mathbf{b}$ in 2 steps

Triangular $\mathbf{Lc} = \mathbf{b}$ Triangular $\mathbf{Ux} = \mathbf{c}$

Then $\mathbf{Ax} = \mathbf{LUx} = \mathbf{Lc} = \mathbf{b}$

$$A = \begin{bmatrix} 1 & 1 & 3 \\ 2 & 3 & 7 \\ 1 & 4 & 6 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 2 & 3 \\ 1 & 4 \end{bmatrix} \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 1 \end{bmatrix} = CR$$

Columns 1 and 2 of A are independent: rank 2

Those independent columns go into C

Column 3 of $A = 2(\text{column 1}) + 1(\text{column 2})$

So 2 and 1 go into column 3 of $R = [I \ F]$

The $n - r$ dependent columns of A
are combinations CF of the r independent columns

Every $A = CR = C \begin{bmatrix} I & F \end{bmatrix} P =$
 $\begin{bmatrix} \text{Independent cols} & \text{Dependent cols} \end{bmatrix}$ Permute cols

Suppose the r columns of C are columns j_1, \dots, j_r of A

Then P puts the r columns of I into columns j_1, \dots, j_r of R

$$\begin{bmatrix} 1 & 2 & 3 & 4 \\ 1 & 2 & 4 & 5 \end{bmatrix} = \begin{bmatrix} 1 & 3 \\ 1 & 4 \end{bmatrix} \begin{bmatrix} 1 & 2 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} = CR$$
$$R = \begin{bmatrix} 1 & 0 & 2 & 1 \\ 0 & 1 & 0 & 1 \end{bmatrix} P$$



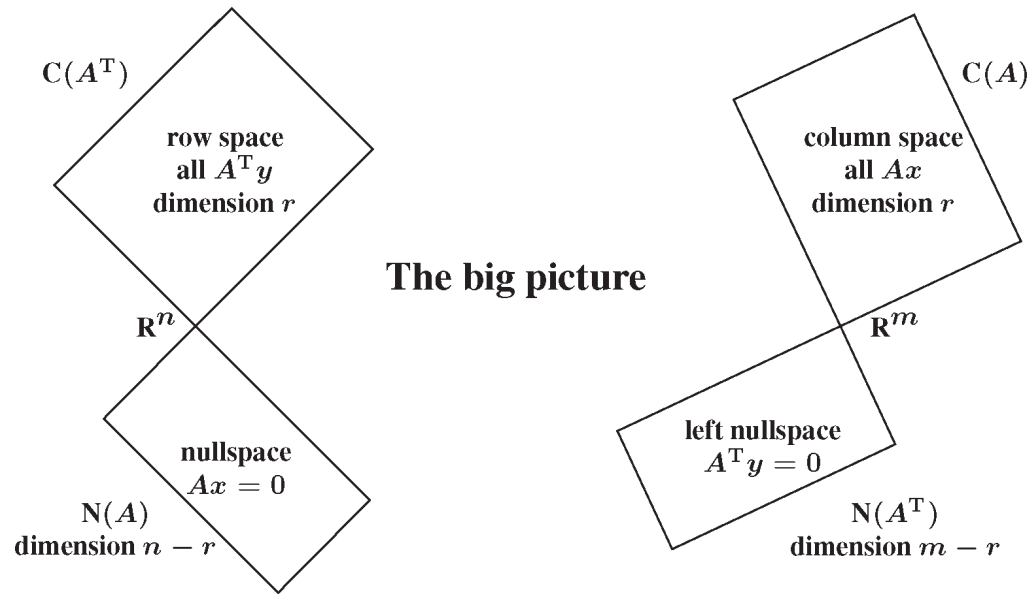
C has r independent columns // R has r independent rows

From $A = CR$, columns of C span the column space of A

From $A = CR$, the rows of R span the row space of A

Column space and row space of A have same dimension r

The Big Picture : Four Subspaces



Elimination computes the reduced row echelon form of A

$$\mathbf{rref}(A) = \begin{bmatrix} R \\ 0 \end{bmatrix} = \begin{bmatrix} I & F \\ 0 & 0 \end{bmatrix} P \quad \begin{array}{l} \text{Finding } R \text{ first} \\ I \text{ locates indep. columns} \end{array}$$

First $k + 1$ columns: $\begin{bmatrix} I_k & F_k \\ 0 & 0 \end{bmatrix} P_k$ are followed by $\begin{bmatrix} \mathbf{u} \\ \boldsymbol{\ell} \end{bmatrix}$

If $\boldsymbol{\ell} = \mathbf{0}$ then \mathbf{u} joins F_k to produce F_{k+1}

If $\boldsymbol{\ell} \neq \mathbf{0}$ then $\begin{bmatrix} \mathbf{u} \\ \boldsymbol{\ell} \end{bmatrix}$ becomes $\begin{bmatrix} \mathbf{0} \\ 1 \\ \mathbf{0} \end{bmatrix}$ to produce I_{k+1}

Note: Every part I, F, P of $\mathbf{rref}(A)$ is determined by A

$$\begin{aligned}
 x_1 + 2x_2 + 11x_3 + 17x_4 = 0 &\longrightarrow x_1 + 3x_3 + 5x_4 = 0 \\
 3x_1 + 7x_2 + 37x_3 + 57x_4 = 0 &\longrightarrow x_2 + 4x_3 + 6x_4 = 0 \\
 4x_1 + 9x_2 + 48x_3 + 74x_4 = 0 &\longrightarrow 0 = 0 \\
 \mathbf{Ax} = \mathbf{0} &\longrightarrow \mathbf{Rx} = \mathbf{0}
 \end{aligned}$$

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 11 & 17 \\ 3 & 7 & 37 & 57 \\ 4 & 9 & 48 & 74 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 3 & 7 \\ 4 & 9 \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 & 5 \\ 0 & 1 & 4 & 6 \end{bmatrix} = \mathbf{CR}$$

Nullspace basis $\begin{bmatrix} -3 \\ -4 \\ 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} -5 \\ -6 \\ 0 \\ 1 \end{bmatrix}$ = Solve $\mathbf{Rx} = \mathbf{0}$ to find
 = “special solutions to $\mathbf{Ax} = \mathbf{0}$ ”

**Block
elimination**

$$P_r A P_c = \begin{bmatrix} W & H \\ J & K \end{bmatrix} \rightarrow \begin{bmatrix} I & W^{-1}H \\ 0 & 0 \end{bmatrix}$$

The “intersection” of r independent rows of A
with r independent columns of A
produces an r by r invertible matrix W

Elimination reduces W to I and finds $R = \begin{bmatrix} I & W^{-1}H \end{bmatrix}$

Factorization 3 “Gram-Schmidt”

$$\mathbf{A} = \mathbf{QR} = (\text{orthogonal})(\text{triangular})$$

From independent columns in \mathbf{A} to

orthogonal unit vectors \mathbf{q}_1 to \mathbf{q}_n

$$\mathbf{Q}^T \mathbf{Q} = \begin{bmatrix} \cdot & \mathbf{q}_1^T & \cdot \\ \cdot & \cdots & \cdot \\ \cdot & \mathbf{q}_n^T & \cdot \end{bmatrix} \begin{bmatrix} \mathbf{q}_1 & \cdots & \mathbf{q}_n \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ \cdot & \cdot & \cdot \\ 0 & 0 & 1 \end{bmatrix} = \mathbf{I}$$

Why is \mathbf{Q} so good? Least squares finds $\hat{\mathbf{x}}$ to minimize $\|\mathbf{Ax} - \mathbf{b}\|^2$

$\mathbf{A}^T \mathbf{A} \hat{\mathbf{x}} = \mathbf{A}^T \mathbf{b} \rightarrow \mathbf{R}^T \mathbf{Q}^T \mathbf{Q} \mathbf{R} \hat{\mathbf{x}} = \mathbf{R}^T \mathbf{Q}^T \mathbf{b} \rightarrow \mathbf{R} \hat{\mathbf{x}} = \mathbf{Q}^T \mathbf{b}$ is easy to solve

Eigenvalues	$S\mathbf{x} = \lambda\mathbf{x}$	Singular values	$A\mathbf{v} = \sigma\mathbf{u}$
$S = S^T$	Orthogonal eigenvectors \mathbf{x}	Orthogonal	$\mathbf{v}_1, \dots, \mathbf{v}_r$
		Orthogonal	$\mathbf{u}_1, \dots, \mathbf{u}_r$

$$S \begin{bmatrix} \mathbf{x}_1 \\ \dots \\ \mathbf{x}_n \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1 \\ \dots \\ \mathbf{x}_n \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \quad A \begin{bmatrix} \mathbf{v}_1 \\ \dots \\ \mathbf{v}_r \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 \\ \dots \\ \mathbf{u}_r \end{bmatrix} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix}$$

$$SX = X\Lambda \quad S = X\Lambda X^{-1} = X\Lambda X^T \quad AV = U\Sigma \quad A = U\Sigma V^T$$

Every matrix $A = (\text{rotation } U) (\text{stretching } \Sigma) (\text{rotation } V^T)$
 (\mathbf{v} 's are eigenvectors of $A^T A$) (\mathbf{u} 's are eigenvectors of AA^T)




Figure: Angle θ from AV to AW is below 90° . Angle $180 - \theta$ from AW to $-AV$ is above 90° . Somewhere in between, as v moves from V toward W , the angle from Av to Aw is exactly 90° . The pictures don't show vector lengths.

Interpolation / Approximation using Deep Learning?

Given: Values w_k at N points u_k in \mathbf{R}^n

Goal: Learning function $F(x, u)$ with weights x and
 $F(x, u_k) \approx w_k$

Chain of functions

$$F(x, u) = F_L(x_L, F_{L-1}(x_{L-1}, F_{L-2}(\dots F_1(x_1, u))))$$

Composition of L simple functions $F_k: \text{vector} \rightarrow \text{vector}$

The weights $\boldsymbol{x} = x_L, x_{L-1}, \dots, x_1$ are chosen so that $F(\boldsymbol{x}, u_k) \approx w_k$ Each $x_k =$ matrix A_k and vector b_k

$$F_k(x_k, F_{k-1}) = \text{ReLU}(A_k F_{k-1} + b_k)$$

Nonlinear $\text{ReLU}(y) = \begin{cases} y & y \geq 0 \\ 0 & y \leq 0 \end{cases}$ applied to each component

Compute weights x to minimize the loss $\|F(x, u) - w\|$

Gradient descent or Stochastic gradient descent : not Newton

MIT OpenCourseWare
<https://ocw.mit.edu>

Resource: A Vision of Linear Algebra
Gilbert Strang

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.