How to Reveal the Rank of a Matrix?

Anil Damle[1], Silke Glas[2], Alex Townsend[1], Annan Yu[1]
[1]Cornell University [2]University of Twente

What is a Rank Revealer of a matrix?

Every matrix $A \in \mathbb{R}^{m\times n}$ admits a singular value decomposition (SVD). The $(k+1)$st singular value of A measures the proximity of A to its nearest rank-k matrix in the spectral norm.

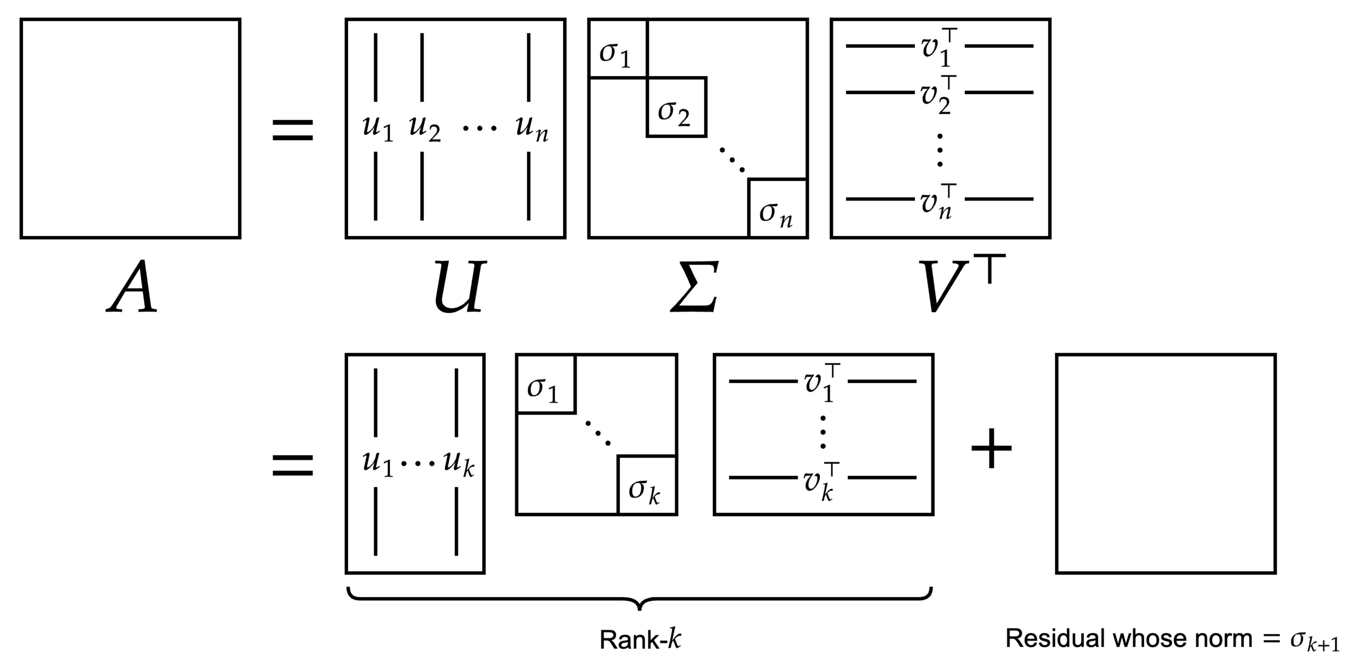

Let $A_k \in \mathbb{R}^{m\times n}$ be a rank-k matrix. We say that A_k is a rank-revealer of A if A_k well-approximates the k leading singular values of A and $A - A_k$ well-approximates the $(\min(m,n)-k)$ trailing singular values of A.

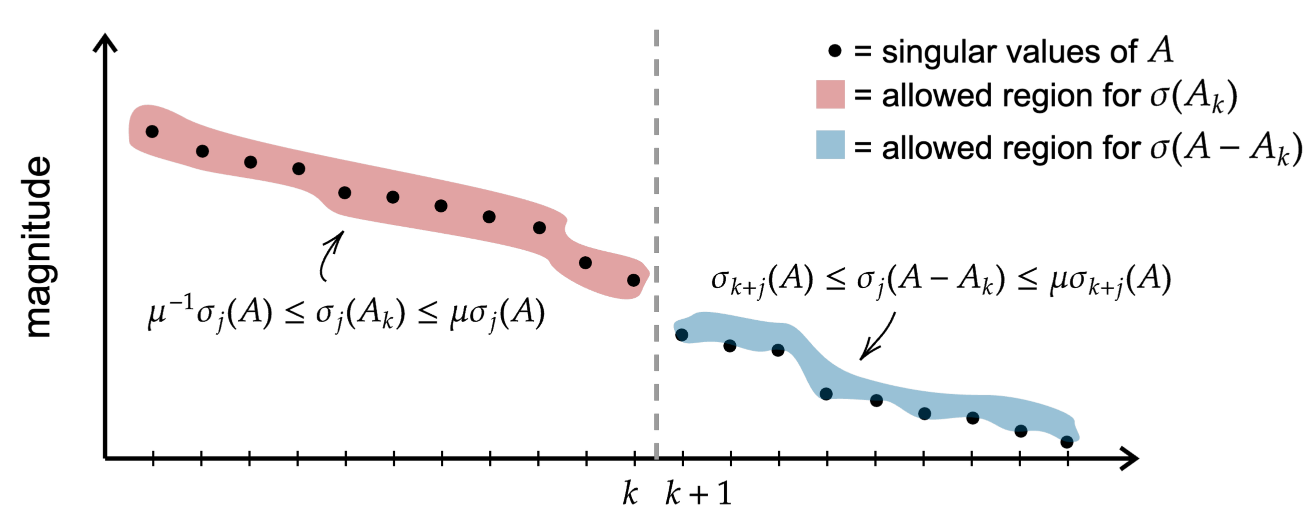

What is the volume of a matrix?

The volume of a matrix is the product of its singular values. Let $B \in \mathbb{R}^{p\times q}$ be a submatrix of A. Associated with it is a quantity γ that measures "how much we can increase its volume by swapping at most one row and at most one column:"

$$\text{vol}(B) \geq \gamma^{-1}\text{vol}(\hat{B})$$

for any submatrix $\hat{B}$ of A that differs in at most one row and column from B. If $\gamma = 1$, then B has a local maximum volume. Local maximum volume does *not* imply global maximum volume.

$$A = \begin{bmatrix} 1 & 3 & 0 & 0 \\ 3 & 1 & 0 & 0 \\ 0 & 0 & \sqrt{3} & 2 \\ 0 & 0 & 2 & -\sqrt{3} \end{bmatrix}$$

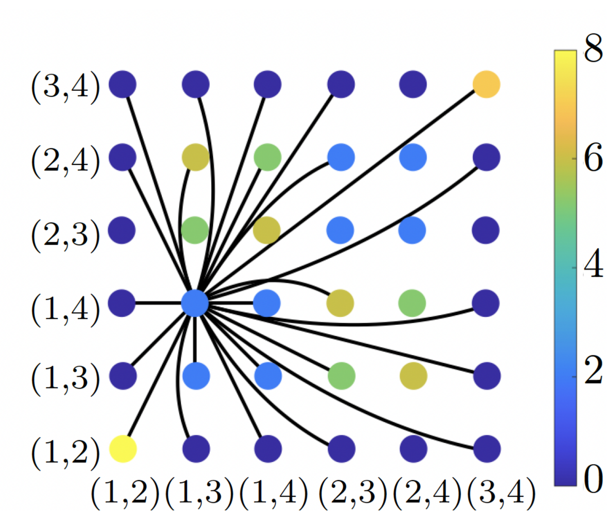

Rank-revealing LU Factorizations

In computing the LU factorization (i.e., Gaussian elimination), a key concept to avoid numerical instabilities is pivoting. That is, we permute the rows and columns of A based on some metric. Some common pivoting strategies include partial, rook, and complete pivoting.

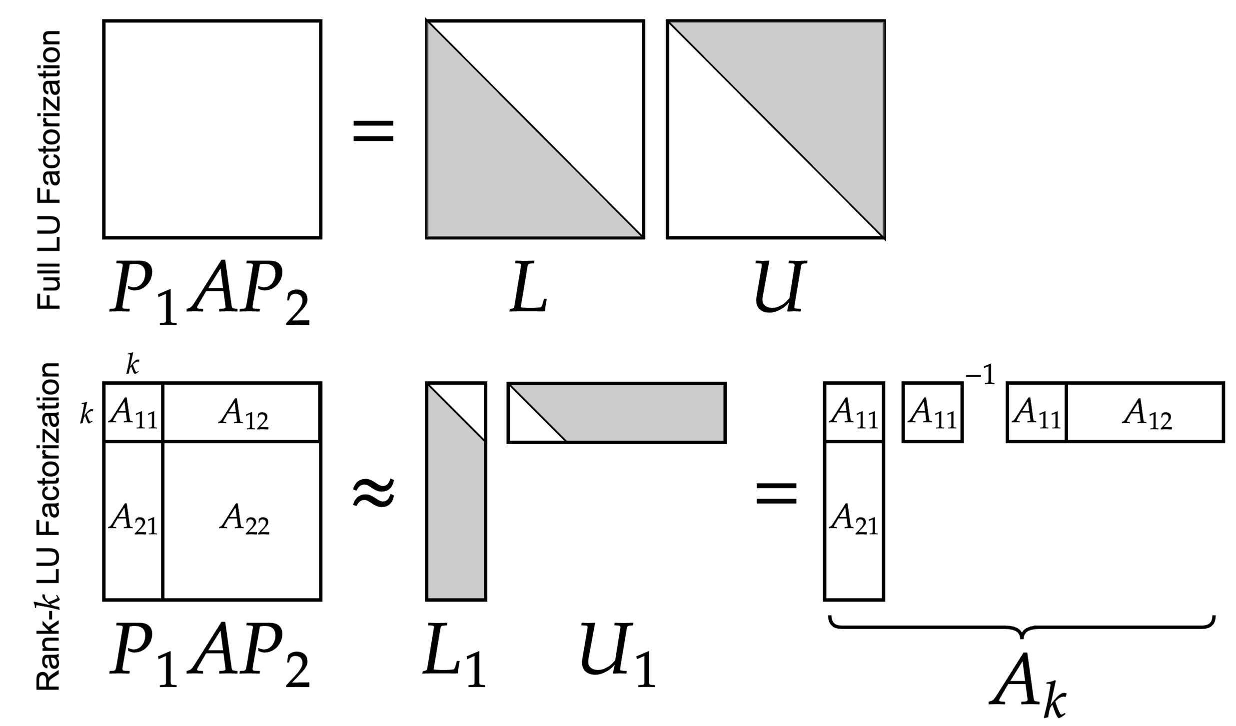

Instead of doing a full LU factorization $P_1AP_2 = LU$, we consider a partial LU factorization, which computes a rank-k approximation A_k of A. The matrix A_k has the nice property that its row (resp. column) space is spanned by the rows (resp. columns) of A. We show that

A_{11} is near local maximum volume $\Leftrightarrow$ A_k is rank-revealing.

Hence, the only pivoting strategy that gives us a rank-revealing approximation is one that guarantees the local maximum volume.

Rank-revealing QR Factorizations

Similarly, for QR factorizations, a pivoting strategy permutes the columns of A. We use a partial QR factorization to construct a rank-k approximation A_k of A, where the column space of A_k is spanned by the columns of A.

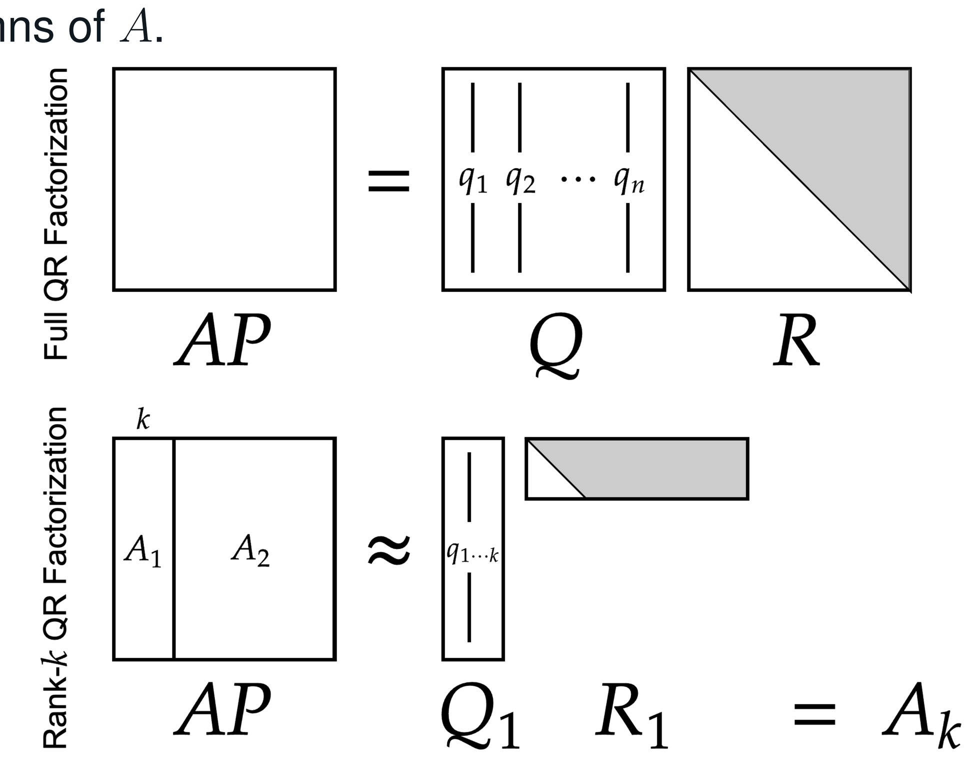

We show that local maximum volume is necessary and sufficient for a rank-revealing QR factorization.

A_1 is near local maximum volume $\Leftrightarrow$ A_k is rank-revealing.

Local Maximum Volume in Practice

Rank-revealing LU and QR factorizations can be computed efficiently in practice. The overall idea is to use an efficient pivoting strategy first, which guarantees a weak lower bound on the pivot volume, and then swap rows and columns to increase it.

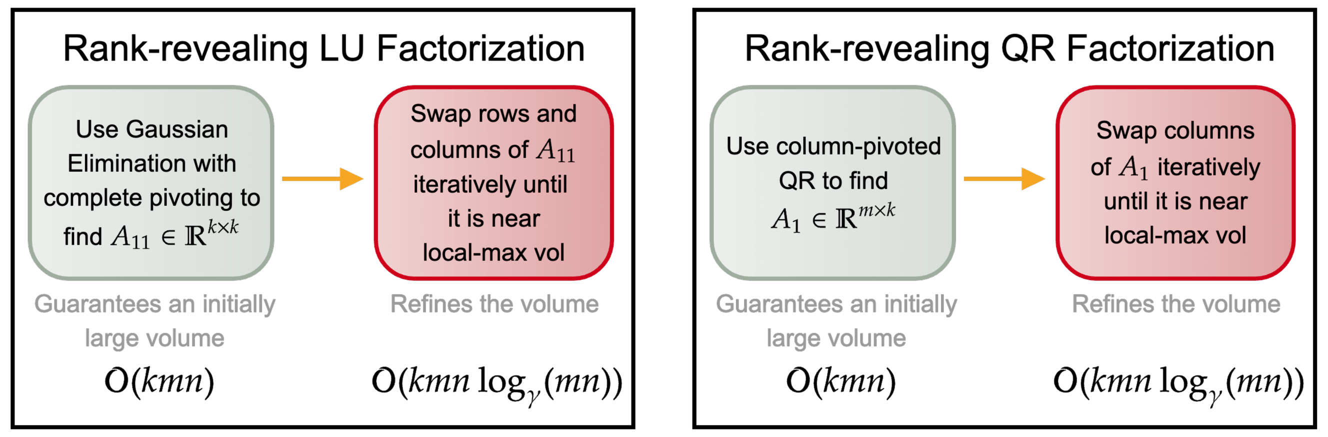

Although the worst-case γ's for GECP and CPQR are large, they are empirically observed to pivot on submatrices with near-local maximum volume, and very few subsequent swaps are needed.

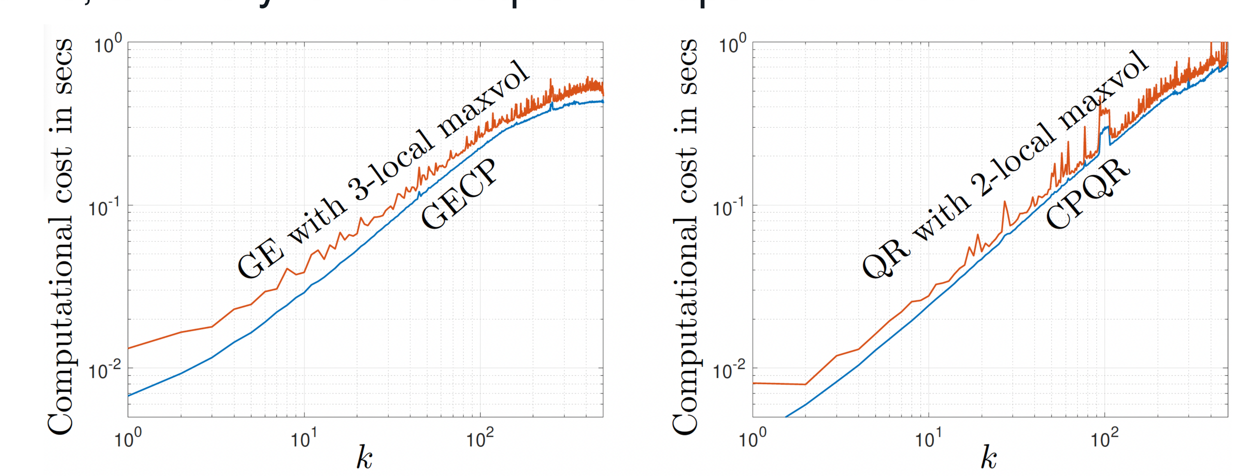

Applications

GECP can be used to approximate functions on a rectangular domain (continuous analog of matrices). We compute γ and observe that GECP generally selects pivots with a large volume. Hence, GECP is empirically a rank-revealer in pseudoskeleton approximation of kernels.

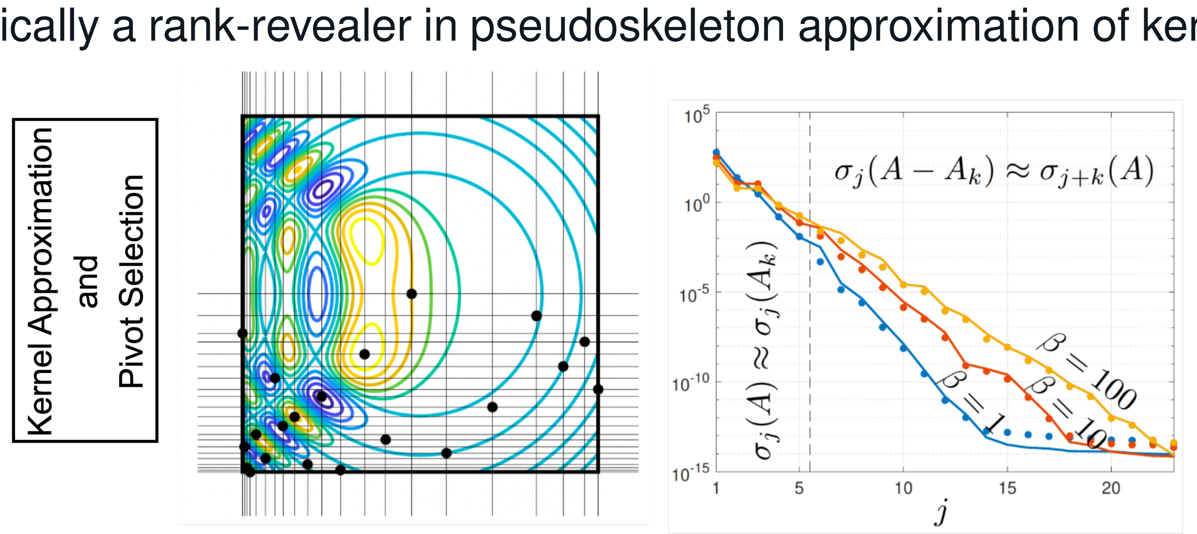

CPQR is used in quantum chemistry to identify a well-conditioned basis as candidates for localized Wannier functions. The columns it selects have a very large local volume.

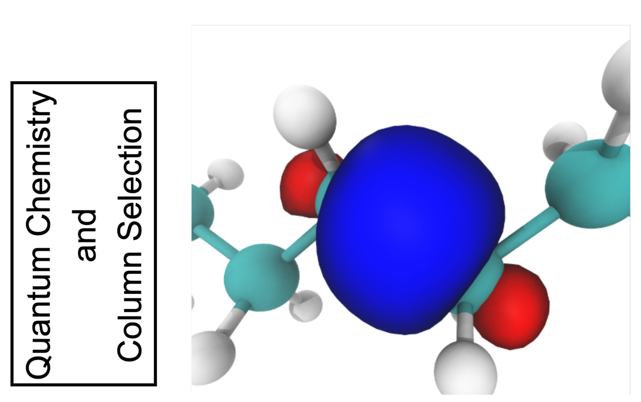

# Columns (n)	# Basis Functions (k)	γ from CPQR	# Swaps to get to $\gamma = 1$
820,125	100	1.0012	10
1,953,125	256	1.0211	9