COL726 Assignment 3

20 March – 3 April, 2021

Note: All answers should be accompanied by a rigorous justification, unless the question explicitly states that a justification is not necessary.

Updated text is highlighted in blue.

1. Define a "PET matrix" to be a square matrix of the form $\mathbf{PTQ}$, where $\mathbf{P}, \mathbf{Q}$ are permutation matrices and $\mathbf{T}$ is triangular.

   (a) Show that there is no need to specify whether $\mathbf{T}$ is upper triangular or lower triangular in the above definition. That is, if $\mathbf{A} = \mathbf{PUQ}$ where $\mathbf{U}$ is upper triangular, there also exist permutation matrices $\mathbf{P}', \mathbf{Q}'$ such that $\mathbf{A} = \mathbf{P}'\mathbf{L}\mathbf{Q}'$ with lower triangular $\mathbf{L}$.

   (b) In the general case, determining if a given matrix $\mathbf{A}$ is a PET matrix is NP-hard. However, suppose we know $\mathbf{A} = \mathbf{PLQ}$ where $\mathbf{L}$ has no extra zeros, i.e. $l_{ij} \neq 0$ for all $i \geq j$. For this case, give an $O(m^2)$ algorithm for finding $\mathbf{P}, \mathbf{Q}$.

2. (a) The standard Cholesky factorization $\mathbf{A} = \mathbf{LL}^*$ requires the use of square roots, which can be expensive on some hardware. Give a modified algorithm without square roots, which computes a factorization $\mathbf{A} = \mathbf{LDL}^*$ where $\mathbf{L}$ is lower triangular and $\mathbf{D}$ is diagonal.

   (b) The modified algorithm no longer requires that the top-left entry a_{11} is positive. Does this mean that the algorithm now works for all nonsingular Hermitian matrices, whether positive definite or not? Prove or give a counterexample.

3. In this question, we investigate why direct methods fail to preserve sparsity. Let a c-sparse matrix be a matrix with at most c nonzero entries in each row and each column.

   (a) Show that the product of two c-sparse matrices is only c^2-sparse.

   (b) Show that for any m, there exists a 3-sparse H.P.D. matrix $\mathbf{A} \in \mathbb{C}^{m\times m}$ whose Cholesky factor is not at all sparse; in particular it has a row or column with $\geq m/2$ nonzero entries.

4. Suppose we attempt to do GMRES without the Arnoldi iteration. At the nth iteration, we simply add a column to the Krylov matrix $\mathbf{K}_n = \begin{bmatrix} \mathbf{b} & \mathbf{Ab} & \cdots & \mathbf{A}^{n-1}\mathbf{b} \end{bmatrix}$ and assume $\mathbf{x} = \mathbf{K}_n\mathbf{y}$ in the least-squares solve.

   (a) Give the full details of the algorithm and analyze its operation count, assuming that computing $\mathbf{A}^{n-1}\mathbf{b}$ takes $\sim cm$ flops. Bonus marks if you can compute the QR factorization for least squares in $O(mn)$ time by reusing the factorization from the previous iteration.

(b) The real problem with this approach is that the conditioning of the least-squares problem grows exponentially with n. Show that this is true even for Hermitian matrices, $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^*$. You may consider the special case $\mathbf{b} = \mathbf{v}_1 + \cdots + \mathbf{v}_m$ and $\|\mathbf{A}\|_2 = \max_{j=1,\ldots,m} |\lambda_j| \neq 1$.

Hint: Note that for any matrix $\mathbf{M} \in \mathbb{C}^{m\times n}$ with $m \geq n$, we have $\kappa(\mathbf{M}) = \dfrac{\sup_{\|\mathbf{x}\|=1} \|\mathbf{M}\mathbf{x}\|_2}{\inf_{\|\mathbf{x}\|=1} \|\mathbf{M}\mathbf{x}\|_2}$.

5. Suppose you apply the conjugate gradient method to a matrix $\mathbf{A}$ whose eigenvalues lie in n disjoint clusters, $[a_1, b_1], [a_2, b_2], \ldots, [a_n, b_n]$, with $\lambda_{\min} = a_1 < b_1 < a_2 < b_2 < \cdots < a_n < b_n = \lambda_{\max}$. Derive a bound on the $\mathbf{A}$-norm of the error after $n, 2n, 3n, \ldots$ CG iterations. Your bound should be tighter than Theorem 38.5; in particular, if all $a_i = b_i$, it should predict that the error becomes 0 after n iterations.

6. (a) Implement a Python function `(P, L, U) = lup(A)` to perform LU factorization with partial pivoting, and a function `x = solveLup(P, L, U, b)` to solve $\mathbf{A}\mathbf{x} = \mathbf{b}$ using the results of `lup`. Following Trefethen & Bau's Exercise 22.2, find a way to construct a matrix $\mathbf{A}$ for which catastrophic rounding errors occur, and implement a function `A = instabilityMatrix(m)` which returns such a matrix of size $m \times m$.

   You should implement backsubstitution yourself instead of using Scipy's built-in function (`scipy.linalg.solve_triangular`).

   (b) Suppose at each step, you have some way of choosing a pivot x_{ij} anywhere in the remaining submatrix, not just in the current column. Give the full algorithm, analogous to Algorithm 21.1, to compute an LU factorization with pivoting in this case.

   (c) A practical strategy for choosing the pivot is "rook pivoting": Find the largest (in absolute value) entry in the current column, and go to its row; find the largest entry in that row, and go to its column; repeat until you have found an entry that is the largest in both its row and its column. Using rook pivoting, implement your algorithm in (b) as a function `(P, Q, L, U) = lupq(A)`, with a corresponding solve function `x = solveLupq(P, Q, L, U, b)`.

   Note: You may use the fact that any permutation matrix is orthogonal.

   For each $m = 1, 2, \ldots, 60$, compute the growth factor ρ (TB 22) of your `instabilityMatrix` for both partial pivoting and rook pivoting, and plot them as a function of m with a logarithmic y-axis. Do the same for the relative backward error in solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ for a randomly chosen $\mathbf{b} \in \mathbb{R}^m$.

Collaboration policy: Refer to the policy on the course webpage.

If you collaborated with others to solve any question(s) of this assignment, give their names in your submission. If you found part of a solution using some online resource, give its URL.

Submission: Submit a PDF of your answers for all questions to Gradescope. Submit the code for Question 6 to Moodle. Both submissions must be uploaded before the assignment deadline.

Code submissions should contain a single .py file which contains all the functions. Functions are permitted but not required to produce any side-effects like printing out values or drawing plots. Any results you are asked to show should go in the PDF.