Homework set 2 — MATH 397C — Spring 2022

Due on Thursday March 3. Please hand in solutions to problems 1, 3, 6, and 7.

Problem 1: Let $\mathbf{A}$ be an $n \times n$ symmetric positive definite matrix. You know that $\kappa(\mathbf{A}) = 100$.

    (1) Suppose that you run conjugate gradients (CG) to solve a linear system $\mathbf{A}\mathbf{x} = \mathbf{y}$. In the worst case scenario, how many iterations would you need in order to *guarantee* that the initial residual has been reduced by a factor of $\varepsilon = 10^{-5}$?

    (2) Suppose that $\mathbf{C}$ is a non-singular matrix and set $\mathbf{B} = \mathbf{C}^{-*}\mathbf{A}\mathbf{C}^{-1}$. You know that $\|\mathbf{I} - \mathbf{B}\| \leq 1/2$. If you use $\mathbf{C}$ as a symmetric pre-conditioner, then how would your answer to question (a) change?

Notes: Equation (38.10) in Trefethen & Bau may be helpful. In part (b), a matrix such as $\mathbf{C}$ can come from an approximate Cholesky factorization $\mathbf{A} \approx \mathbf{C}^*\mathbf{C}$. In such a case, $\mathbf{C}$ is upper triangular (and often very sparse), which means that applying either $\mathbf{C}^{-*}$ or $\mathbf{C}^{-1}$ to a vector is inexpensive.

Problem 2: Let $\mathbf{A}$ be an $m \times n$ matrix, set $p = \min(m, n)$, and suppose that the singular value decomposition of $\mathbf{A}$ takes the form

(1)
$$\underset{m \times n}{\mathbf{A}} = \underset{m \times p}{\mathbf{U}} \quad \underset{p \times p}{\mathbf{D}} \quad \underset{p \times n}{\mathbf{V}^*}.$$

Let k be an integer such that $1 \leq k < p$ and let $\mathbf{A}_k$ denote the truncation of the SVD to the first k terms:

$$\mathbf{A}_k = \mathbf{U}(:, 1:k)\,\mathbf{D}(1:k, 1:k)\,\mathbf{V}(:, 1:k)^*.$$

Recall the definitions of the spectral and Frobenius norms:

$$\|\mathbf{A}\| = \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|}, \qquad \text{and} \qquad \|\mathbf{A}\|_{\mathrm{F}} = \left(\sum_{i=1}^{m} \sum_{j=1}^{n} |\mathbf{A}(i, j)|^2 \right)^{1/2}.$$

Prove directly from the definitions of the norms that

$$\|\mathbf{A} - \mathbf{A}_k\| = \sigma_{k+1}$$

and that

$$\|\mathbf{A} - \mathbf{A}_k\|_{\mathrm{F}} = \left(\sum_{j=k+1}^{p} \sigma_j^2 \right)^{1/2}.$$

Problem 3: On the course webpage, download the file `hw02p3.m`. This file contains an implementation of the column pivoted QR algorithm that computes a rank-k approximation to a matrix, for any given k. Your task is now to do two modifications to the code:

(a) Starting with the function `CPQR_given_rank` write a new function with calling sequence

$$[Q,R,ind] = CPQR\_given\_tolerance(A,acc)$$

that takes as input an accuracy, and computes a low-rank approximation to a matrix that is accurate to precision "acc".

(b) Write a function with calling sequence

$$[U,D,V] = SVD\_given\_tolerance(A,acc)$$

that computes a diagonal matrix **D**, and orthonormal matrices **U** and **V** such that

$$\|\mathbf{A} - \mathbf{UDV}^*\| \leq \varepsilon,$$

where ε is the given tolerance. The idea is to use the function `CPQR_given_tolerance(A,acc)` that you created in part (a).

Please hand in a print-out of the code that you created.

Extra problem: The file `hw02p3.m` creates a plot that shows the accuracies of two low-rank approximations: The truncated SVD on the one hand, and the truncated QR on the other. Let me encourage you to play around with this a bit, try different matrices and see how the approximations errors compare. There is no need to hand anything in!

Problem 4: In this example, we investigate the effect *blocking* has on execution time of matrix computations.

(a) Suppose that we are given two $n \times n$ matrices **B** and **C** and that we seek to compute $\mathbf{A} = \mathbf{BC}$. We could do this in Matlab either by just typing `A = B*C`, or, we could write a loop

for $i = 1:n$

    `A(:,i) = B*C(:,i)`

end for

The code `hw02p4.m` illustrates the two techniques. It turns out that while the two methods are mathematically equivalent, doing it via a loop is much slower. In this problem, please measure the time T_n required for several different values of n. Test the hypothesis that $T_n = C n^3$ by plotting your measure valued of T_n versus n in a log-log-diagram. Fit a straight line through the points, and estimate C. Hand in the graph and the values of C that you estimate for the two methods.

(b) Repeat the problem in (a), but now compare three different matrix factorization algorithms:
  - `[Q,R] = qr(A)`
    This factorization can be blocked. It is fast, but no good for low-rank approximation.
  - `[Q,R,J] = qr(A,'vector')`
    Column pivoted QR factorization — intermediately fast, and good for low-rank approximation.
  - `[U,D,V] = svd(A)`
    Singular value decomposition — slowest, but excellent for low-rank approximation.

Problem 5: Recall the *single pass* RSVD shown in Figure 1. Consider the following piece of Matlab code:

```
m=12; n=20; b=3; k=5;
A   = rand(m,n);
Gc  = randn(n,k);
Gr  = randn(m,k);
Yc  = zeros(m,k);
Yr  = zeros(n,k);

for i = 1:(n/b)
  ind        = (i-1)*b + (1:b);
  A_slice    = A(:,ind);
  Yc         = Yc + A_slice*Gc(ind,:);
  ?????????????????????????????????????
end

fprintf(1,'Error in Yc = %12.3e\n',max(max(abs(Yc - A *Gc))))
fprintf(1,'Error in Yr = %12.3e\n',max(max(abs(Yr - A'*Gr))))
```

This snippet of code emulates how a streaming algorithm would interact with $\mathbf{A}$ — you read a set of b columns at a time, and use the information in $\mathbf{A}_{\text{slice}}$ to build $\mathbf{Y}_{\text{c}}$ piece by piece.

(a) Write code to replace the question marks. The result of the new code should be that after the loop completes, the matrix `Yr` has also been computed. Note that this line should reference only `A_slice`, not `A` itself. (This can be solved by a single line of code, but if you use more, then that is fine too.)

(b) Currently, the code only works if `n` is an integer multiple of `b`. Modify the code so that it works for any block size `b`.

Hand in the code you write.

ALGORITHM: SINGLE-PASS RANDOMIZED SVD FOR A GENERAL MATRIX

Inputs: An $m \times n$ matrix $\mathbf{A}$, a target rank k, and an over-sampling parameter p (say $p = 10$).

Outputs: Matrices $\mathbf{U}$, $\mathbf{V}$, and $\mathbf{D}$ in an approximate rank-k SVD of $\mathbf{A}$. (I.e. $\mathbf{U}$ and $\mathbf{V}$ are ON and $\mathbf{D}$ is diagonal.)

Stage A:

(1) Form two Gaussian random matrices $\mathbf{G}_{\text{c}} = \text{randn}(n, k + p)$ and $\mathbf{G}_{\text{r}} = \text{randn}(m, k + p)$.
(2) Form the sample matrices $\mathbf{Y}_{\text{c}} = \mathbf{A}\,\mathbf{G}_{\text{c}}$ and $\mathbf{Y}_{\text{r}} = \mathbf{A}^{*}\,\mathbf{G}_{\text{r}}$.
(3) Form ON matrices $\mathbf{Q}_{\text{c}}$ and $\mathbf{Q}_{\text{r}}$ consisting of the k dominant left singular vectors of $\mathbf{Y}_{\text{c}}$ and $\mathbf{Y}_{\text{r}}$.

Stage B:

(4) Let $\mathbf{C}$ denote the $k \times k$ least squares solution of the joint system of equations formed by the equations $\left(\mathbf{G}_{\text{r}}^{*}\mathbf{Q}_{\text{c}}\right)\mathbf{C} = \mathbf{Y}_{\text{r}}^{*}\mathbf{Q}_{\text{r}}$ and $\mathbf{C}\left(\mathbf{Q}_{\text{r}}^{*}\mathbf{G}_{\text{c}}\right) = \mathbf{Q}_{\text{c}}^{*}\mathbf{Y}_{\text{c}}$.
(5) Decompose the matrix $\mathbf{C}$ in a singular value decomposition $[\hat{\mathbf{U}}, \mathbf{D}, \hat{\mathbf{V}}] = \text{svd}(\mathbf{C})$.
(6) Form $\mathbf{U} = \mathbf{Q}_{\text{c}}\hat{\mathbf{U}}$ and $\mathbf{V} = \mathbf{Q}_{\text{r}}\hat{\mathbf{V}}$.

FIGURE 1. A randomized single-pass algorithm suitable for a general matrix.

Problem 6: On the course webpage, download the file `hw02p6.m`. This file contains an implementation of the basic RSVD scheme. It computes approximate matrix factorizations using the basic RSVD, and plots the approximation error versus the minimum error as produced by the truncated (full) SVD. The error is reported in both the spectral and the Frobenius norms.

For this problem, code up the single pass algorithm described in Figure 1 and include it in the comparison. Hand in a printout of your implementation of the single pass algorithm (you do not need to print out the driver code, etc, just the actual subroutine). Also hand in the error plots for three different sets of test matrices. In the code that you can download, two test cases are already included. You are welcome to use these two. For the third, come up with some matrix you find interesting yourself! It could be dense, sparse, etc. Just remember that for any of this to make sense, the singular values of the matrix you pick must show at least some degree of decay.

Note: In step (4) of the algorithm, the matrix **C** is determined by jointly solving two matrix equations. Note that this is a bit complicated to implement. For simplicity, simply pick one of the two equations and determine **C** by solving that one, ignoring the other.

Problem 7: The purpose of this exercise is to investigate how well certain structured random embeddings do at approximating the range of a matrix. Given an $m \times n$ matrix **A**, we seek to build a sample from its range of the form $\mathbf{Y} = \mathbf{A\Omega}$, where $\mathbf{\Omega}$ is a random matrix of size $n \times \ell$. We consider three different distributions:

(1) A standard Gaussian random matrix.
(2) A subsampled randomized discrete Fourier transform, $\mathbf{\Omega} = \mathbf{DFS}$, where **D** is a diagonal matrix whose entries are drawn randomly from $\{-1, 1\}$, where **F** is the discrete cosine transform (or DFT, if you prefer), and where **S** is an $n \times \ell$ sampling matrix formed by randomly drawing a subset of ℓ columns from the $n \times n$ identity matrix.
(3) A sparse random matrix formed by randomly placing ζ nonzero entries in each row of $\mathbf{\Omega}$, where ζ is a small integer (say $\zeta \in \{2, 3, 4, \ldots, 8\}$). Each non-zero entry is drawn from a standard Gaussian distribution.

Code up these choices, and write a script that evaluates the error

$$e_\ell = \|\mathbf{A} - \mathbf{P_Y A}\|,$$

where $\mathbf{P_Y}$ is the orthogonal projection onto col(**Y**), as a function of the number of samples ℓ drawn. Hand in some graphs that show how the error e_ℓ compares to the theoretically minimal value $\sigma_{\ell+1}$ for a few different matrices **A**. For instance, you can use the three test matrices provided in the routine `hw02p6.m`. However, add at least one additional matrix of your choice. Briefly describe the matrices that you use.

Hint 1: There is no need to code the SRFT up efficiently. I would simply apply the full FFT or the full DFT to create a big sample matrix $\mathbf{Y}_{\text{big}} = \mathbf{ADF}$, then draw ℓ columns at random from this matrix to form **Y**.

Hint 2: Note that $\mathbf{P_Y} = \mathbf{QQ}^*$ where **Q** is the tall thin "Q" factor in a QR factorization of **Y**.

Background: The preprints https://arxiv.org/abs/2104.05877 and https://arxiv.org/abs/2002.01387 may be of interest to anyone wanting to learn more about structured random embeddings. There is no need to study these in order to solve this problem, however!