Exercise set 8 – MATH-250 Advanced Numerical Analysis I

The exercise sheet is divided into two sections: quiz and exercises. The quiz will be discussed in the beginning of the lecture on Thursday, April 17. The exercises marked with ($\star$) are graded homework. The exercises marked with **(Python)** are implementation based and can be solved in the Jupyter notebooks which are available on Moodle/Noto. **The deadline for submitting your solutions to the homework is Friday, May 2 at 10h15.**

Quiz

a) If A is not invertible then A does not have an LU factorization (without pivoting).

☐ True ☐ False

b) If A is invertible then Algorithm 4.13 in the lecture notes does not fail, that is, it always finds nonzero pivot elements $|a_{ik}|$ (in exact arithmetic) and produces an LU factorization with pivoting for A.

☐ True ☐ False

c) The norm defined by $\|A\|_{\max} := \max_{ij} |a_{ij}|$ is a matrix norm but it is not submultiplactive.

☐ True ☐ False

d) On the vector space of square symmetric matrices, $\operatorname{trace}(A) = a_{11} + \cdots + a_{nn}$ is a matrix norm.

☐ True ☐ False

e) Given a diagonal matrix $A = \operatorname{diag}(a_{11}, \ldots, a_{nn})$, which of the following statements is *wrong*?

☐ $\|A\|_F = \|\mathbf{d}\|_2$ for the vector $\mathbf{d} = [a_{11}, a_{22}, \ldots, a_{nn}]$

☐ $\|A\|_2 = \|\mathbf{d}\|_\infty$

☐ $\|A\|_1 = \|\mathbf{d}\|_1$

Exercises

Problem 1.

(a) Show that for $\mathbf{x} \in \mathbb{R}^n$

(i) $\|\mathbf{x}\|_2 \le \|\mathbf{x}\|_1 \le \sqrt{n}\|\mathbf{x}\|_2$

(ii) $\|\mathbf{x}\|_\infty \le \|\mathbf{x}\|_1 \le n\|\mathbf{x}\|_\infty$

(iii) $\|\mathbf{x}\|_\infty \le \|\mathbf{x}\|_2 \le \sqrt{n}\|\mathbf{x}\|_\infty$

In addition, show that the bounds are tight.

(b) For $A \in \mathbb{R}^{n\times n}$ and $p \ge 1$, the matrix p-norm of A is defined as

$$\|A\|_p = \sup_{\mathbf{x}\neq 0} \frac{\|A\mathbf{x}\|_p}{\|\mathbf{x}\|_p}$$

Show that

(i) $\frac{1}{\sqrt{n}}\|A\|_2 \le \|A\|_1 \le \sqrt{n}\|A\|_2$

(ii) $\frac{1}{\sqrt{n}}\|A\|_\infty \le \|A\|_2 \le \sqrt{n}\|A\|_\infty$

(c) For $A \in \mathbb{R}^{n\times n}$ the Frobenius norm of A is defined as

$$\|A\|_F = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}^2}$$

Show that $\|A\|_F = \sqrt{\mathrm{tr}(A^TA)} = \sqrt{\mathrm{tr}(AA^T)}$ where $\mathrm{tr}(A) = \sum_{i=1}^{n} a_{ii}$.

(d) Show that for $A, B \in \mathbb{R}^{n\times n}$ and $p \ge 1$ we have

(i) $\|AB\|_p \le \|A\|_p\|B\|_p$ using the definition of the matrix p-norm.

(ii) $\|AB\|_F \le \|A\|_F\|B\|_F$

(e) Show that for $A \in \mathbb{R}^{n\times n}$

(i) $\|A\|_1 = \max_{1\le j\le n} \sum_{i=1}^{n} |a_{ij}|$

(ii) $\|A\|_\infty = \max_{1\le i\le n} \sum_{j=1}^{n} |a_{ij}|$

Problem 2.

Consider the matrix $A \in \mathbb{R}^{10\times 10}$ and the vector $\mathbf{b} \in \mathbb{R}^{10}$ given below

$$A = \begin{pmatrix} 1 & 10 & & & 0 \\ & 1 & 10 & & \\ & & \ddots & \ddots & \\ & & & 1 & 10 \\ 0 & & & & 1 \end{pmatrix} \quad \mathbf{b} = \begin{pmatrix} 9 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

(a) Solve $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$ exactly.

(b) Solve the perturbed system $A\widehat{\mathbf{x}}^{(1)} = \mathbf{b} + \Delta\mathbf{b}$ where $\Delta\mathbf{b} = 10^{-8}\mathbf{e}_3$, where

$$\mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

(c) Compute the relative error

$$\frac{\|\mathbf{x} - \widehat{\mathbf{x}}^{(1)}\|_\infty}{\|\mathbf{x}\|_\infty}$$

(d) Now consider the perturbed system $(A + \Delta A)\widehat{\mathbf{x}}^{(2)} = \mathbf{b}$ where $\Delta A = \varepsilon I_{10}$. Using results about sensitivity of linear systems, what is the maximum value of ε so that the relative error

$$\frac{\|\mathbf{x} - \widehat{\mathbf{x}}^{(2)}\|_\infty}{\|\mathbf{x}\|_\infty}$$

is guaranteed to be less than 10^{-3}? Use Python to compute `np.linalg.cond(A, np.inf)` and/or `np.linalg.norm(np.linalg.inv(A), np.inf)`, if necessary.

(e) Let $\varepsilon = 10^{-6}$ and solve the system $(A + \Delta A)\widehat{\mathbf{x}}^{(2)} = \mathbf{b}$ in Python and compute the relative error.

Problem 3.

A matrix A is strictly diagonally dominant by rows if

$$|a_{ii}| > \sum_{\substack{j=1 \\ j \neq i}}^{n} |a_{ij}|, \quad i = 1, \dots, n.$$

By using the Neumann series that you have seen in Proposition 4.22, show that a strictly diagonally dominant matrix is non-singular.

Hint: Without loss of generality, assume that the diagonal entries of A all equal 1 by a suitable scaling of the rows of A. Now recall what you know about some matrix norms, such as the operator norms induced by $\|\cdot\|_1$ and $\|\cdot\|_\infty$.

($\star$) **Problem 4.** (Do note the later submission deadline due to the Easter break.)

Let $k\colon [0,1] \times [0,1] \to \mathbb{R}$ and $u\colon [0,1] \to \mathbb{R}$ be continuous functions. We define the integral operator $F\colon [0,1] \to \mathbb{R}$ as the integral of u with the kernel k using

$$F(x) = \int_0^1 k(x,y)u(y)\,\mathrm{d}y. \tag{1}$$

For a partition of $[0,1]$ into $N > 0$ subintervals denote $h = \frac{1}{N}$ and let Q_h be the composite trapezoidal rule on the N subintervals of length h. Further define the subinterval's boundary points $x_i = i \cdot h$ for $i = 0, 1, \dots, N$.

(a) We want to apply Q_h to approximate the operator (1) at each x_i

$$Q_h[k(x_i,\cdot)u(\cdot)] = \hat{F}(x_i) \approx F(x_i) = \int_0^1 k(x_i,y)u(y)\,\mathrm{d}y, \quad i = 0, 1, \dots, N.$$

To this end we define the function value vectors

$$\hat{\mathbf{f}} = [\hat{F}(x_0), \hat{F}(x_1), \dots, \hat{F}(x_N)]^\top \quad \text{and} \quad \mathbf{u} = [u(x_0), u(x_1), \dots, u(x_N)]^\top.$$

Show that there exists an $(N+1) \times (N+1)$ matrix A such that $A\mathbf{u} = \hat{\mathbf{f}}$. Provide explicit formulae for the entries a_{ij} of A.

(b) We now consider the opposite idea of (a). Given a vector of the integral operator's evaluations $\mathbf{f} = [F(x_0), F(x_1), \ldots, F(x_N)]^\top$ we solve $A\hat{\mathbf{u}} = \mathbf{f}$, and use the result to approximate $F(z)$ for any arbitrary value of $z \in [0,1]$.

Assume that the matrix

$$K = \begin{pmatrix} k(x_0,x_0) & k(x_0,x_1) & \cdots & k(x_0,x_N) \\ k(x_1,x_0) & k(x_1,x_1) & \cdots & k(x_1,x_N) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_N,x_0) & k(x_N,x_1) & \cdots & k(x_N,x_N) \end{pmatrix} \tag{2}$$

is invertible and show that

$$\hat{F}(z) = [k(z,x_0), k(z,x_1), \ldots, k(z,x_N)]K^{-1}\mathbf{f} \tag{3}$$

holds true.

(c) For $N = 4$ suppose that the corresponding matrix K from (b) is invertible.

Show that $\hat{F}(x_i) = F(x_i)$ for $i = 0,1,2,3,4$.

(d) We choose the *radial basis function kernel* $k(x,y) = \exp\left(-(x-y)^2/4\right)$.

Implement a Python function `approximate_operator(F, N, z)` that computes the vector $\hat{F}(\mathbf{z}) = [\hat{F}(z_1), \hat{F}(z_2), \ldots, \hat{F}(z_m)]^\top$ given a vector $\mathbf{z} = [z_1, z_2, \ldots, z_m]^\top, m > 0$, using (3). Assure that your implementation requires $\mathcal{O}(N^3 + mN^2)$ operations.

(e) Let $N \in \{2, 5, 10\}, m = 1000, \mathbf{z} =$ `np.linspace(0, 1, num=1000)`, and define

$$F_1(x) = \sin(3\pi x) \quad \text{and} \quad F_2(x) = \exp\left(-|x - 0.5|^{2/3}\right).$$

For each N and each F_i plot the true function $F_i(\mathbf{z})$ and its approximation $\hat{F}_i(\mathbf{z})$. Compute the maximum absolute error $\max_{j=1,2,\ldots,m} |F_i(z_j) - \hat{F}_i(z_j)|$ and clearly display this error.

(f) Explain the behaviour for the approximation of F_2 with $N = 10$.

To mitigate this bad approximation we utilize regularisation. This means that instead of (3) we compute

$$\hat{F}^{(\gamma)}(z) = [k(z,x_0), k(z,x_1), \ldots, k(z,x_N)](K + \gamma\ \mathrm{id})^{-1}\mathbf{f}$$

for some small $\gamma > 0$.

Implement a Python function `approximate_operator_reg(F, N, z, gamma)` to compute $\hat{F}^{(\gamma)}$ similarly to (d); you may reuse your code from (d). Determine a value for γ such that the maximum absolute error of the approximation for F_2 and $N = 10$ is less than or equal to 10^{-1}.

(g) **Bonus (This part is not needed to get full marks.)**: Prove that the matrix K from (2) is always symmetric and positive semidefinite for the radial basis kernel $k(x, y) = \exp(-(x-y)^2/4)$. You can use the Schur product theorem or any other technique.

Schur Product Theorem. Let $A, B \in \mathbb{R}^{n \times n}$ be two symmetric and positive semidefinite matrices. Then their elementwise product $(A \odot B)_{ij} = a_{ij} \cdot b_{ij}$ is once again symmetric and positive semidefinite.

Remember to upload a scan `homework08.pdf` of your solutions and the completed Jupyter notebook `homework08.ipynb` corresponding to the homework to the submission panel on Moodle until Friday, May 2 at 10h15. To download your notebook from Noto, use `File > Download`. Only your submissions to Moodle will be considered for grading.