

# ALOP Autumn School

## High performance solution of linear optimization problems

### Numerical linear algebra

1. Load the matrix  $A$  from the file `bcsstk08.mat`: this defines an array  $A$

- Identify the dimension  $n$  of the matrix  $A$
- Generate  $\mathbf{xstar}$  to represent a vector  $\mathbf{x}^* \in \mathbb{R}^n$ , all of whose entries are 1
- Compute  $\mathbf{b}=\mathbf{A}*\mathbf{xstar}$  to represent the vector  $\mathbf{b} = A\mathbf{x}^*$
- Find and report the time required to solve  $A\mathbf{x} = \mathbf{b}$ 
  - By computing  $A^{-1}$  as `Ainv` and then forming `Ainv*b`
  - Using `A\b`

You should see that it is much faster to use “backslash” than form the inverse and multiply by it.

- Compute and report the residual norm  $\|A\mathbf{x} - \mathbf{b}\|_2$  and norm of the solution error  $\|\mathbf{x} - \mathbf{x}^*\|_2$ . Is one approach more accurate than the other?
- Repeat using the files `nasa1824.mat` and `pds02.mat`. Do your observations change?

2. Download `Arrowmatrix.m` which, for a given value of  $n$ , returns  $A =$

$$A = \begin{bmatrix} -1 & -1 & -1 & \dots & -1 & -1 \\ -1 & 1 & & & & \\ -1 & & 1 & & & \\ \vdots & & & \ddots & & \\ -1 & & & & 1 & \\ -1 & & & & & 1 \end{bmatrix}$$

- Convince yourself that forming the decomposition  $PA = LU$  using Gaussian elimination with partial pivoting results in  $L$  and  $U$  having no zero entries below and above the diagonal respectively. Convince yourself that forming the decomposition  $PAQ = LU$  by pivoting using the Markowitz criterion leads to  $L$  and  $U$  having no more nonzero entries than the corresponding triangles of  $A$ .
- Write a script (using the dimension `n=10` to develop your script) which performs the following
  - Determine the number of nonzeros in  $A$  [using `nnz(A)`]
  - Determine the elapsed time to compute  $A^{-1}$  [using `inv(A)`]
  - Determine the total number of nonzeros in  $A^{-1}$
  - Determine the elapsed time to form the decomposition  $PA = LU$  [using `[L U P]=lu(A)`]
  - Determine the total number of nonzeros in both  $L$  and  $U$  [ie `nnz(L)+nnz(U)-n`, see below for explanation of the subtraction of  $n$ ]
  - Report the dimension of  $A$ , its number of nonzeros, the times for the matrix computations in (ii) and (iv), and the numbers of nonzeros in their results
- The command `sparse(A)` returns the matrix  $A$  in a format which exploits its sparsity. Add commands to your script to perform the following
  - Form the sparse version of  $A$ , referred to below as `SparseA`
  - Determine the elapsed time to form the (sparsity-exploiting) decomposition  $PAQ = LU$  [using `[L U P Q]=lu(SparseA)`]
  - Determine the total number of nonzeros in both  $L$  and  $U$

Run your script for  $n = 50 \times 2^k$  for  $k = 1, \dots, 7$ . You should observe that it is much more efficient (in terms of time and sparsity) to form  $PAQ = LU$  than  $PA = LU$  or  $A^{-1}$ .

- (d) After a computation with a sparse matrix, the ratio of the number of nonzeros in the result to the number of nonzeros in the matrix is known as the *fill factor*. The best that can reasonably be expected is a fill factor of 1. The density of the result of a matrix computation is the number of nonzeros in the result divided by the product of the dimensions of the matrix, and is usually expressed as a percentage.

Add commands to your script to compute and report the fill factor and density for the computed LU decompositions and inverse

Note

- Since the  $n$  diagonal entries of  $L$  are known structurally,  $n$  is subtracted from the total number of nonzeros in  $L$  and  $U$  to get the number of true nonzeros. Otherwise the (trivial) LU decomposition of the identity matrix would have a fill factor of 2.
- To force `sprintf` to display a “%” symbol, you need to put “%%” into the string. This is because a single “%” symbol is interpreted as a format statement.

- (e) Download `bcsstk08.mat`, `nasa2910.mat` and `c-36.mat`, and run your script for the matrix  $A$  which is loaded by the command `load('file.mat')`; for each of the files `bcsstk08.mat`, `nasa2910.mat` and `c-36.mat`, observing the extreme density and fill-factor when forming  $A^{-1}$  and the relative efficiency of exploiting sparsity when forming the LU decomposition.

- (f) Run your script for the file `pds02.mat`. Do you observe any odd behaviour?

3. Download `MyCG.m` which performs (at most  $n$ ) iterations of the conjugate gradient method, starting from  $\mathbf{x}^{(1)} = \mathbf{0}$ , until  $\|\mathbf{r}^{(k+1)}\|_2 \leq 10^{-8}$ . As well as  $\mathbf{x}$ , the function returns the index  $m$  of the final iteration.

Conjugate gradient algorithm:

For  $A\mathbf{x} = \mathbf{b}$ , given some estimate  $\mathbf{x}^{(1)}$  of  $\mathbf{x}$ , let  $\mathbf{s}^{(1)} = \mathbf{r}^{(1)} = \mathbf{b} - A\mathbf{x}^{(1)}$ . Then, for some tolerance  $\epsilon$ , repeat for  $k = 1, 2, \dots$

$$\mathbf{w} = A\mathbf{s}^{(k)}$$

$$\alpha^{(k)} = \frac{\mathbf{r}^{(k)T} \mathbf{s}^{(k)}}{\mathbf{w}^T \mathbf{s}^{(k)}}$$

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha^{(k)} \mathbf{s}^{(k)}$$

$$\mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} - \alpha^{(k)} \mathbf{w}$$

If  $\|\mathbf{r}^{(k+1)}\|_2 \leq \epsilon$  then stop

$$\beta^{(k)} = \frac{\|\mathbf{r}^{(k+1)}\|_2^2}{\|\mathbf{r}^{(k)}\|_2^2}$$

$$\mathbf{s}^{(k+1)} = \mathbf{r}^{(k+1)} + \beta^{(k)} \mathbf{s}^{(k)}$$

- (a) The performance of the conjugate gradient method depends on the eigenvalues of the matrix  $A$ . Download `AnalyseEigs.m` and extend your script by calling `AnalyseEigs(A)` before performing CG. `AnalyseEigs(A)` counts and reports the number of complex and non-positive eigenvalues of  $A$ , as well as plotting the eigenvalues on a log scale.

Test the performance of CG using the following matrices, choosing the solution of  $A\mathbf{x} = \mathbf{b}$  to be a random vector and forming  $\mathbf{b} = A\mathbf{x}$ .

- $A = \lambda I$  of dimension  $n = 1000$ . Choose any nonzero value of  $\lambda$  that you want: CG should always terminate after one iteration
- The matrix  $A$  from `662_bus.mat`. [Remember to extract `n=size(A,1)`.]
- The matrix  $A$  from `bcsstk08.mat`. [Remember to extract `n=size(A,1)`.]

- iv. The matrix  $A = I + \epsilon E$  of dimension  $n = 1000$ , where  $E$  is the random sparse symmetric matrix with  $\tau = 10n$  entries generated by the MATLAB function `SpSymmRand(n, tau)` which you will have to download. Use the sparse identity matrix given by `speye(n)`. Try  $\epsilon = 10^{-k}$  for  $k = 1, 2, 4, 8$ .
  - v.  $A = I + \mathbf{u}\mathbf{u}^T$  of dimension  $n = 1000$  for a random vector  $\mathbf{u}$  such that  $\|\mathbf{u}\|_2 = 1$ . This matrix has two distinct eigenvalues:  $\lambda = 2$  (of multiplicity 1) and  $\lambda = 1$  (of multiplicity  $n - 1$ ).
- (b) Compute and report the elapsed time of the CG iterations, together with the error measures  $\|\hat{\mathbf{x}} - \mathbf{x}\|_2$  and  $\|A\hat{\mathbf{x}} - \mathbf{b}\|_2$ , where  $\hat{\mathbf{x}}$  is the solution computed by CG. Compare these with the same performance measures for  $\mathbf{A} \setminus \mathbf{b}$ .
- (c) **Extension:** A preconditioner for  $A$  is the **incomplete Cholesky decomposition**  $A = \tilde{L}\tilde{L}^T$ , where  $\tilde{L}$  is computed with values below a “drop tolerance” in magnitude set to zero. In MATLAB this is computed as
- ```
L = ichol(A,struct('type','ict','droptol',droptol));
```
- i. Develop a function `MyPCG.m` which uses  $\tilde{L}$  as a preconditioner. Test this with a small problem such as `662_bus.mat`. As the drop tolerance reduces, so should the number of CG iterations required.
  - ii. For the larger problem `bodyy4.mat`, can you find a value for the drop tolerance such that the time required to compute  $\tilde{L}$  and use it in `MyPCG.m` is less than the time required by MATLAB’s backslash operator?

## Machine learning and data science applications

1. In the **support vector machine** problem, we are given a training dataset of  $n$  points of the form

$$(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n),$$

where  $y_i \in \{-1, 1\}$  indicates the class to which the point  $\mathbf{x}_i \in \mathbb{R}^p$  belongs. Let the rows of  $X \in \mathbb{R}^{n \times p}$  be the vectors  $\mathbf{x}_i$  and let  $\mathbf{y} \in \{-1, 1\}^n$  be the vector of class indicators  $y_i$ . A linear support vector machine classifies a new point  $\mathbf{x} \in \mathbb{R}^p$  by

$$\mathbf{x} \mapsto \text{sign}(\mathbf{w}^T \mathbf{x} + w_0).$$

The normal vector  $\mathbf{w} \in \mathbb{R}^p$  and the offset  $w_0/\|\mathbf{w}\|$  define a hyperplane in  $\mathbb{R}^p$  and need to be determined from the training set. To this end, we need to solve

$$\underset{\mathbf{w}, w_0}{\text{minimize}} \quad \frac{C}{n} \left[ \sum_{i=1}^n \max(0, 1 - y_i(\mathbf{w}^T \mathbf{x}_i - w_0)) \right] + \|\mathbf{w}\|_2^2,$$

where  $C > 0$  is a penalty parameter.

- (a) By introducing variables

$$s_i = \max(0, 1 - y_i(\mathbf{w}^T \mathbf{x}_i - w_0)) \quad i = 1, \dots, n$$

show that the problem may be formulated as the smooth, constrained QP problem

$$\underset{\mathbf{s}, \mathbf{w}, w_0}{\text{minimize}} \quad \mathbf{w}^T I \mathbf{w} + \frac{C}{n} \mathbf{e}^T \mathbf{s} \quad \text{subject to} \quad -YX\mathbf{w} + \mathbf{y}w_0 - \mathbf{s} \leq -\mathbf{e}, \quad \mathbf{s} \geq \mathbf{0},$$

where  $\mathbf{e}$  is a vector of ones,  $Y$  is a diagonal matrix and  $\mathbf{y}$  is a vector, both with entries  $y_i$  for  $i = 1, \dots, n$ , and  $s_i$  is the  $i^{\text{th}}$  component of  $\mathbf{s}$ .

- (b) Load `svm_training.mat` into your MATLAB session, which contains a set of 1000 training points in  $\mathbb{R}^2$  and their class indicators. Visualize the points using the provided `svmplot.m` function.

- (c) Solve the quadratic program for this data set and penalty parameter  $C = 100$ . The MATLAB function `quadprog` can be used to

$$\text{minimize } \frac{1}{2} \mathbf{x}^T H \mathbf{x} + \mathbf{c}^T \mathbf{x} \quad \text{subject to } A \mathbf{x} \leq \mathbf{b}$$

[**Note:** To pass a vector of lower bounds on the variables, the function is called with `x=quadprog(H,c,A,b,[],[],lb)`.]

- (d) Add the hyperplane [line since  $p = 2$ ] to the plot of points created by `svmplot.m`

2. The ball  $B = \{\mathbf{x} \mid \|\mathbf{y} - \mathbf{x}\|_2 \leq r\}$  of centre  $\mathbf{y} \in \mathbb{R}^n$  and radius  $r$  is the set of all points in  $\mathbb{R}^n$  within a Euclidean distance  $r$  of  $\mathbf{y}$ . The **Chebyshev centre problem** is that of finding the centre and radius of the largest ball within the polyhedron  $P = \{\mathbf{x} \in \mathbb{R}^n \mid A \mathbf{x} \leq \mathbf{b}\}$ , where  $A \in \mathbb{R}^{m \times n}$ .

- (a) i. Show that points a distance  $r$  from  $\mathbf{y} \in \mathbb{R}^n$  in the direction  $\mathbf{a}_i$  are given by

$$\mathbf{x} = \mathbf{y} + \frac{r}{\|\mathbf{a}_i\|_2} \mathbf{a}_i$$

- ii. Show that these points lie within the set  $\mathbf{a}_i^T \mathbf{x} \leq b_i$  if

$$b_i \geq \mathbf{a}_i^T \mathbf{y} + \|\mathbf{a}_i\|_2 r$$

- iii. Hence show that the Chebyshev centre problem may be formulated as the LP problem

$$\text{minimize}_{\mathbf{y}, r} -r \quad \text{subject to } \mathbf{a}_i^T \mathbf{y} + \|\mathbf{a}_i\|_2 r \leq b_i, \quad i = 1, \dots, m$$

- (b) Find the Chebyshev centre and maximum radius when

$$A = \begin{bmatrix} 1 & 1 & 1 \\ 2 & 0 & 1 \\ 0 & -1 & -1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 15 \\ 12 \\ -5 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

[**Note:** Calling `x=linprog(c,A,b)` in MATLAB finds the solution of the LP problem

$$\text{minimize } \mathbf{c}^T \mathbf{x} \quad \text{subject to } A \mathbf{x} \leq \mathbf{b},$$

where arrays  $A$ ,  $\mathbf{b}$ ,  $\mathbf{c}$  and  $\mathbf{x}$  correspond to a matrix  $A$  and vectors  $\mathbf{b}$ ,  $\mathbf{c}$  and  $\mathbf{x}$ .]

- (c) Which of the hyperplanes does the maximum ball touch?

[Note that you should use a tolerance on values being zero (e. g.  $10^{-6}$ ) since the LP solution in (b) is computed using inexact arithmetic.]

- (d) How does a small change in the last component of  $\mathbf{b}$  affect the radius?

[Adapted from D. Michaels, K. Ballerstein ETH Zürich, spring term 2011]

## Interior point methods

### 1. IPM Exercise:

This is an opportunity to use a simple interior point solver written in MATLAB to solve small LP problems

(a) Download `ipexample.m`, `ipdriver.m`, `newtonsolve.m` and `newtonsolver.m`

- The function `ipdriver.m` solves the LP problem

$$\text{minimize } f = \mathbf{c}^T \mathbf{x} \quad \text{s. t.} \quad \mathbf{A}\mathbf{x} = \mathbf{b} \quad \text{and} \quad \mathbf{x} \geq \mathbf{0}$$

- The script `ipexample.m` sets up  $\mathbf{A}$ ,  $\mathbf{b}$  and  $\mathbf{c}$  and calls `ipdriver.m` to solve

$$\begin{aligned} \min f &= 2x_1 + x_2 \\ \text{s. t.} \quad &x_1 + x_2 \geq 3 \\ &3x_1 + x_2 \geq 6 \\ &x_1 \geq 0 \quad x_2 \geq 0 \end{aligned}$$

(b) Modify `ipexample.m` so that it solves

$$\begin{aligned} \max f &= x_1 + 2x_2 + 3x_3 \\ \text{s. t.} \quad &x_1 + 2x_3 \leq 3 \\ &x_2 + 2x_3 \leq 2 \\ &x_1 \geq 0 \quad x_2 \geq 0 \quad x_3 \geq 0 \end{aligned}$$

(c) To use `ipdriver.m` to solve larger LP problems, the values of  $\mathbf{A}$ ,  $\mathbf{b}$  and  $\mathbf{c}$  can be read from `.mat` files. Download `AFIRO.mat` and `ADLITTLE.mat` then use `ipdriver.m` to solve them.

(d) **Extension:** Study `ipdriver.m`, `newtonsolve.m` and `newtonsolver.m` to see how the IPM is implemented in MATLAB

(e) **Extension:** Modify `ipdriver.m` and `newtonsolver.m` so that the IPM implementation solves QP problems.