Norwegian University of Science and Technology
Department of Mathematical Sciences

TMA4180
Optimisation I
Spring 2020

Solutions to exercise set 4

1 Assume that $A \in \mathbb{R}^{m\times n}$ is a matrix and that $b \in \mathbb{R}^m$.

a) Show that $x^* \in \mathbb{R}^n$ solves the *least squares problem*

$$\min_{x\in\mathbb{R}^n} \|Ax - b\|^2, \tag{1}$$

if and only if x^* satisfies the *normal equations*

$$A^\top A x^* = A^\top b.$$

Solution: The least squares problem is an unconstrained minimisation problem for the function $f(x) = \|Ax-b\|^2$ on $\mathbb{R}^n$. Observe that f is smooth, and that

$$\nabla f(x) = 2A^\top(Ax-b) \qquad \text{and} \qquad \nabla^2 f(x) = 2A^\top A.$$

Calculation of ∇f follows either from the chain rule in the multivariable setting, or by direct expansion

$$\|Ax-b\|^2 = (Ax-b)^\top(Ax-b) = x^\top A^\top A x - 2b^\top Ax + b^\top b.$$

The matrix $A^\top A$ is symmetric, and also positive semi-definite, because

$$v^\top A^\top A v = (Av)^\top Av = \|Av\|^2 \geq 0 \quad \text{for all} \quad v \in \mathbb{R}^n.$$

Hence, f is convex and we infer that every critical point is a global minimiser (and conversely). As such, x^* minimises f if and only if $\nabla f(x^*) = 0$. In other words,

$$A^\top A x^* = A^\top b.$$

b) Show that the optimization problem (1) admits a solution $x^* \in \mathbb{R}^n$.

Solution: There are many ways to show this, but we shall use the previous result. From linear algebra we have that the column space of $A^\top A$ and $A^\top$ coincide,

$$\operatorname{col} A^\top A = \operatorname{col} A^\top.$$

If the reader is not familiar with this result, the reader can prove this by first showing that $\ker A^\top A = \ker A$ (note that $Ax = 0 \Leftrightarrow 0 = \|Ax\|^2 = x^\top A^\top Ax$) and combine this with the matrix identity $(\ker X)^\perp = \operatorname{col} X^\top$. Using this result, the trivial statement $A^\top b \in \operatorname{col} A^\top$ implies that $A^\top b \in \operatorname{col} A^\top A$, and so we deduce the existence of $x^* \in \mathbb{R}^n$ satisfying

$$A^\top A x^* = A^\top b.$$

By the previous exercise, $x^* \in \mathbb{R}^n$ is also a solution of (1).

c) Show that the solution x^* of (1) is unique, if the rank of A equals n.

Solution: If $\operatorname{rank} A = n$ it means that the columns of A are linearly independent, and therefore the homogeneous problem $Av = 0$ admits only a trivial solution. As a result, the Hessian of our objective function is positive definite; indeed

$$v^\top \nabla^2 f(x) v = v^\top A^\top A v = \|Av\|^2 \geq 0$$

with equality only when $v = 0$. Consequently the function is strictly convex, and the global minimum is unique.

d) Show that, regardless of the rank of A, the optimization problem

$$\min_{x\in\mathbb{R}^n} \|x\|^2 \qquad \text{s.t. } x \text{ solves (1)} \tag{2}$$

admits a unique solution $x^\dagger \in \mathbb{R}^n$.

Solution: We have already shown that the function f is convex, and that its set of global minimizers is non-empty regardless of A. Owing to the convexity of f, its set of global minimizers is also a convex set; let us call it Ω — these are precisely the points satisfying (1). Clearly Ω is closed (this is true for any l.s.c. function f). Therefore, a continuous and coercive function $g(x) = \|x\|^2$ admits at least one minimizer on Ω. Further, since g is strictly convex ($\nabla^2 g = 2I$), there cannot be more than one minimizer in Ω (otherwise their convex combination would be an even better solution in Ω).

2 Assume that $A \in \mathbb{R}^{n\times n}$ is symmetric and positive *semi*-definite and $b \in \operatorname{Range} A$. Show that (in exact arithmetics) the CG algorithm converges for every starting point $x_0 \in \mathbb{R}^n$ in at most $m = \dim(\operatorname{Range} A)$ iterations to a solution of $Ax = b$.

(This shows that at least theoretically the assumption of positive definiteness can be slightly relaxed.)

Solution: The main difficulty is in showing that the algorithm does not break down with divisions by zero when the steplength α_k is computed, as there are could be directions $p \neq 0$ such that $p^\top Ap = 0$. For this to be the case, however, the direction p needs to be in $\ker A$: indeed, if we expand $p = \sum_i c_i v_i$ in terms of orthonormal eigenvectors v_i of A, which correspond to eigenvalues $\lambda_i \geq 0$, then $p^\top Ap = \sum_i \lambda_i c_i^2$. For the latter sum to be zero p must be a linear combination of eigenvectors, corresponding to the zero eigenvalue.

We will first show that throughout the usual CG algorithm we maintain $p_k \in \operatorname{Range} A$ so that divisions by zero are avoided. We will then show the estimate on the number of iterations.

At iteration 0 we have $p_0 = r_0 = b - Ax_0 \in \operatorname{Range}(A) - \operatorname{Range}(A) \in \operatorname{Range} A$. Assuming that $p_k \in \operatorname{Range}(A)$, we compute $p_{k+1} = r_{k+1} + \beta_{k+1} p_k \in \operatorname{Range}(A) - \beta_{k+1} \operatorname{Range}(A) \in \operatorname{Range}(A)$, because $r_{k+1} = b - Ax_{k+1} \in \operatorname{Range}(A) - \operatorname{Range}(A) \in \operatorname{Range} A$.

The usual inductive proof of convergence of CG implies that the algorithm constructs orthogonal residuals $\{r_0, r_1, \dots\}$ and conjugate directions $\{p_0, p_1, \dots\}$. Normally we rely on the fact that the number of conjugate or orthogonal directions in the n-dimensional space is n, therefore the algorithm must converge in at most n steps. Hovever, all residuals are by construction in Range A, which in the present case has dimension $m \leq n$. Thus the algorithm will generate a zero residual (in exact arithmetics) after at most m steps.

3 Assume that $m > n$, that $A \in \mathbb{R}^{m \times n}$, and that $b \in \mathbb{R}^m$. Consider the following algorithm:

- Choose $x_0 \in \mathbb{R}^n$ arbitrary, set $r_0 \leftarrow Ax_0 - b$, $s_0 \leftarrow A^\top r_0$, $p_0 \leftarrow -s_0$, and $k \leftarrow 0$.
- While $s_k \neq 0$:

$$\begin{aligned}
\alpha_k &\leftarrow \frac{\|s_k\|^2}{\|Ap_k\|^2}, \\
x_{k+1} &\leftarrow x_k + \alpha_k p_k, \\
r_{k+1} &\leftarrow r_k + \alpha_k A p_k, \\
s_{k+1} &\leftarrow A^\top r_{k+1}, \\
\beta_{k+1} &\leftarrow \frac{\|s_{k+1}\|^2}{\|s_k\|^2}, \\
p_{k+1} &\leftarrow -s_{k+1} + \beta_{k+1} p_k, \\
k &\leftarrow k + 1.
\end{aligned}$$

Assume that the matrix A has full rank. Show that the algorithm above is actually identical with the CG-algorithm for the solution of $A^\top A x = A^\top b$ (in the sense that the iterates x_k of both methods coincide).

Solution: We provide an inductive argument, showing that

$$r_{k-1}^{\text{CG}} = s_{k-1}, \qquad p_{k-1}^{\text{CG}} = p_{k-1}, \qquad \alpha_{k-1}^{\text{CG}} = \alpha_{k-1}, \qquad \text{and} \qquad x_k^{\text{CG}} = x_k$$

for any k, assuming x_0 arbitrary but equal for both methods, with superscript "CG" for the CG-parameters. Remark: CG-algorithm is well-defined because $A^\top A$ is symmetric positive definite (rank $A = n$).

Base case $k = 1$ follows from

$$r_0^{\text{CG}} = (A^\top A)x_0 - A^\top b, \qquad r_0 = Ax_0 - b, \qquad \text{and} \qquad s_0 = A^\top r_0 = r_0^{\text{CG}},$$

so that

$$p_0^{\text{CG}} = -r_0^{\text{CG}} = -s_0 = p_0,$$

and

$$\alpha_0^{\text{CG}} = \frac{\|r_0^{\text{CG}}\|^2}{\left(p_0^{\text{CG}}\right)^\top (A^\top A) p_0^{\text{CG}}} = \frac{\|r_0^{\text{CG}}\|^2}{\|Ap_0^{\text{CG}}\|^2} = \frac{\|s_0\|^2}{\|Ap_0\|^2} = \alpha_0.$$

Therefore

$$x_1^{\text{CG}} = x_0 + \alpha_0^{\text{CG}} p_0 = x_0 + \alpha_0 p_0 = x_1.$$

Suppose next that the induction hypothesis is true for some $k \in \mathbb{Z}_+$. Then

$$\begin{aligned} r_k^{\mathrm{CG}} &= r_{k-1}^{\mathrm{CG}} + \alpha_{k-1}^{\mathrm{CG}} A^\top A p_{k-1}^{\mathrm{CG}} \\ &= s_{k-1} + \alpha_{k-1} A^\top A p_{k-1} \\ &= A^\top \left(r_{k-1} + \alpha_{k-1} A p_{k-1} \right) \\ &= A^\top r_k \\ &= s_k, \end{aligned}$$

$$p_k^{\mathrm{CG}} = -r_k^{\mathrm{CG}} + \frac{\left\| r_k^{\mathrm{CG}} \right\|^2}{\left\| r_{k-1}^{\mathrm{CG}} \right\|^2} p_{k-1}^{\mathrm{CG}} = -s_k + \frac{\| s_k \|^2}{\| s_{k-1} \|^2} p_k = p_k,$$

and

$$\alpha_k^{\mathrm{CG}} = \frac{\left\| r_k^{\mathrm{CG}} \right\|^2}{\left\| A p_k^{\mathrm{CG}} \right\|^2} = \frac{\| s_k \|^2}{\| A p_k \|^2} = \alpha_k,$$

so, most importantly,

$$x_k^{\mathrm{CG}} = x_{k-1}^{\mathrm{CG}} + \alpha_{k-1}^{\mathrm{CG}} p_{k-1}^{\mathrm{CG}} = x_{k-1} + \alpha_{k-1} p_{k-1} = x_k.$$

4 In this exercise, we study the Gauß–Newton method for solving the least-squares problem corresponding to the (overdetermined and inconsistent) system of equations

$$\begin{aligned} x + y &= 1, \\ x - y &= 0, \\ xy &= 2. \end{aligned}$$

To that end, we define

$$\begin{aligned} r_1(x, y) &:= x + y - 1, \\ r_2(x, y) &:= x - y, \\ r_3(x, y) &:= xy - 2, \end{aligned}$$

and

$$f(x, y) := \frac{1}{2} \sum_{j=1}^{3} r_j(x, y)^2.$$

We denote moreover by $J = J(x, y)$ the Jacobian of $r = (r_1, r_2, r_3) \colon \mathbb{R}^2 \to \mathbb{R}^3$.

a) Show that the function f is non-convex, but that it has a unique minimiser (x^*, y^*).

Solution: The gradient and Hessian of f equal

$$\nabla f(x, y) = J^\top r = \begin{bmatrix} 1 & 1 & y \\ 1 & -1 & x \end{bmatrix} \begin{bmatrix} x + y - 1 \\ x - y \\ xy - 2 \end{bmatrix} = \begin{bmatrix} 2(x - y) + xy^2 - 1 \\ 2(y - x) + yx^2 - 1 \end{bmatrix}$$

and

$$\begin{aligned} \nabla^2 f(x, y) &= J^\top J + r_1 \nabla^2 r_1 + r_2 \nabla^2 r_2 + r_3 \nabla^2 r_3 \\ &= \begin{bmatrix} 2 + y^2 & xy \\ xy & 2 + x^2 \end{bmatrix} + 0 + 0 + r_3 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 2 + y^2 & 2(xy - 1) \\ 2(xy - 1) & 2 + x^2 \end{bmatrix}. \end{aligned}$$

Since, for example, $\nabla^2 f(-1,1)$ has eigenvalues -1 and 7, it follows that f is non-convex. However, f does have a unique minimiser: it is smooth and coercive, and thus we infer that there is a global minimiser, which must also be a stationary point. Coercivity can be seen this way: if $f \le C^2/2$ for some $C > 0$, then $r_1^2 \le C$ and $r_2^2 \le C$. If we put $z_1 = x+y$ and $z_2 = x-y$, then these inequalities imply $|z_1| \le C+1$ and $|z_2| \le C$. As a consequence we have the boundedness of $|x| = |(z_1+z_2)/2| \le C+1/2$ and $|y| = |(z_1-z_2)/2| \le C+1/2$. Therefore, if either $|x| \to \infty$ or $|y| \to \infty$, then also $f(x,y) \to \infty$.
The stationary point for f must satisfy the equations

$$xy(x+y) = 2 \qquad \text{and} \qquad xy(x-y) = 4(x-y),$$

which can be seen by adding and subtracting the equations in the system $\nabla f = 0$. If $x \ne y$, then $xy = 4$ from the second equation, so that $y = \frac{1}{2} - x$ from the first. But as $4 = xy = x(\frac{1}{2} - x)$ has complex solutions in x, we reject this case. Therefore $x = y$, which gives solutions $x = y = 1$ from the first equation. Thus the function has only one stationary point, and since the minimum exists and must satisfy the optimality conditions, this is the point of global minimum.

b) Show that the matrix $J^\top J$ required in the Gauß–Newton method is positive definite for all x, y.

Solution: Remember first that any matrix of the form $J^\top J$ is symmetric positive semi-definite (SPSD), which follows from

$$v^\top (J^\top J) v = (Jv)^\top (Jv) = \|Jv\|^2 \ge 0.$$

Moreover, SPSD matrices are characterised by having nonnegative eigenvalues, while a matrix is symmetric positive definite (SPD) if and onlfy if it has strictly positive eigenvalues.
Computing $\det J^\top J = 2(x^2+y^2+2) > 0$, we see that $J^\top J$ is invertible. In particular, all eigenvalues are nonzero, and hence, strictly positive (being nonnegative). Therefore, $J^\top J$ is positive definite.

c) Show that the Gauß–Newton method with Wolfe line search for the minimisation of f converges for all initial values (x_0, y_0) to the unique solution of the non-linear least squares problems.

Solution: We show that $J(x,y)$ satisfies the "full-rank condition"

$$\|J(x,y)v\| \ge \gamma \|v\|$$

for all $(x,y) \in \mathbb{R}^2$, where $\gamma > 0$ is a constant. Theorem 10.1 in N&W then implies that the Gauß–Newton method with Wolfe line search converges for all initial values.
Now,

$$\begin{aligned}\|J(x,y)v\|^2 &= (v_1+v_2)^2 + (v_1-v_2)^2 + (yv_1+xv_2)^2 \\ &\ge 2\left(v_1^2+v_2^2\right) = 2\|v\|^2,\end{aligned}$$

and so we may put $\gamma = \sqrt{2}$ to get the desired inequality.

d) Perform one step of the Gauß–Newton method (without line search) for the solution of this least-squares problem. Use the initial value $(x_0, y_0) = (0, 0)$.

Solution: With $(x_0, y_0) = (0, 0)$, we have

$$J^\top J = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \quad \text{and} \quad J^\top r = \begin{bmatrix} -1 \\ -1 \end{bmatrix}.$$

Solving the linear system $J^\top J p = -J^\top r$ gives $p = (1/2, 1/2)$, so that

$$(x_1, y_1) = (x_0, y_0) + p = (1/2, 1/2).$$