Norwegian University of Science and Technology
Department of Mathematical Sciences

TMA4180 Optimisation I Spring 2020

Solutions to exercise set 3

1 Consider the quadratic function

$$f(x) = \tfrac{1}{2}x^\top Qx - b^\top x,$$

where $Q \in \mathbb{R}^{d\times d}$ is a symmetric and positive definite matrix and $b \in \mathbb{R}^d$.

a) Compute the gradient and the Hessian of the function f, verify that f is strictly convex, and find the unique global minimum of f.

Solution: $\nabla f = Qx - b$ and $\nabla^2 f = Q$ from calculus. Since Q is symmetric positive definite (SPD), it follows that f is strictly convex on $\mathbb{R}^d$, and as such, there is at most one global minimum of f. Furthermore, this global minimum x^* must be a stationary point satisfying $\nabla f(x^*) = 0$. We conclude that $x^* = Q^{-1}b$, since Q is invertible (all eigenvalues of Q are positive, and hence, different from zero).

b) Let $x \in \mathbb{R}^d$, and let $p \in \mathbb{R}^d$ be a direction satisfying the inequality $\nabla f(x)^\top p < 0$. Compute analytically the step length $\alpha_{x,p}$ that solves the (exact) linesearch problem $\min_{\alpha>0} f(x + \alpha p)$.

Solution: First of all, to avoid trivial cases let us note that $p \neq 0$ and $\nabla f(x) = Qx - b \neq 0$ owing to the inequality $\nabla f(x)^\top p < 0$.
Now, let us look at the first order necessary conditions for $\alpha_{x,p}$ to be a minimizer:

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} f(x + \alpha_{x,p}p) = p^\top \nabla f(x + \alpha_{x,p}p) = p^\top [Q(x + \alpha_{x,p}p) - b] = 0,$$

or

$$\alpha_{x,p} = -\frac{p^\top [Qx - b]}{p^\top Qp} > 0,$$

since $p^\top Qp > 0$ owing to Q being positive definite, and $p^\top [Qx-b] = p^\top \nabla f(x) < 0$ by our assumption.
Since $\mathrm{d}^2/\mathrm{d}\alpha^2 f(x + \alpha p) = p^\top Qp > 0$ the linesearch problem is strictly convex, and therefore $\alpha_{x,p}$ is the unique global minimum.

c) Let $x, p \in \mathbb{R}^d$ and $\alpha_{x,p} > 0$ be as in the previous question. Show that the steplength $\alpha_{x,p}$ satisfies the strong Wolfe conditions if and only if $c_1 \leq 1/2$.

Solution: Clearly the strong curvature condition is satisfied because $\nabla f(x + \alpha_{x,p}p)^\top p = \mathrm{d}/\mathrm{d}\alpha f(x+\alpha_{x,p}p) = 0$, thus the "new" slope is 0 and must be smaller than or equal in magnitude than the slope we have started with.

We check the sufficient decrease condition now:

$$f(x+\alpha_{x,p}p)-f(x)=\frac{1}{2}\alpha_{x,p}^2 p^\top Qp+\alpha p^\top(Qx-b)=-\frac{1}{2}\frac{[p^\top(Qx-b)]^2}{p^\top Qp}<0,$$

while

$$c_1\alpha_{x,p}\nabla f(x)^\top p=-c_1\frac{[p^\top(Qx-b)]^2}{p^\top Qp}.$$

Thus the sufficient decrease condition is equivalent to the inequality $c_1 \leq 1/2$.

2 *Cf. the exam 2016.*

A function $f\colon \mathbb{R}^d \to \mathbb{R}$ is called *strongly convex*,[1] if there exists $c>0$ such that the function $x \mapsto f(x)-\frac{c}{2}\|x\|^2$ is convex.

a) Show that a twice differentiable function $f\colon \mathbb{R}^d \to \mathbb{R}$ is strongly convex, if and only if there exists $c>0$ such that

$$p^\top \nabla^2 f(x) p \geq c\|p\|^2$$

for all $p \in \mathbb{R}^d$ and $x \in \mathbb{R}^d$.

Verify that this is equivalent to the condition that all the eigenvalues of all the matrices $\nabla^2 f(x)$, $x \in \mathbb{R}^d$, are larger or equal to c.

Solution: A twice differentiable function g is convex if and only if, $\nabla^2 g(x)$ is positive semi-definite for all x, which is equivalent to $p^\top \nabla^2 g(x) p \geq 0$ for all x and p. Applying this characterization to $g(x)=f(x)-\frac{c}{2}\|x\|^2$ with Hessian $\nabla^2 g(x)=\nabla^2 f(x)-cI$, where $I \in \mathbb{R}^{d\times d}$ is the identity matrix, we get that f is strongly convex, if and only if $\nabla^2 f(x)-cI$ is positive semi-definite for some $c>0$ and all $x \in \mathbb{R}^d$, which in turn is equivalent to

$$p^\top(\nabla^2 f(x)-cI)p \geq 0 \Leftrightarrow p^\top \nabla^2 f(x) p \geq c\|p\|^2$$

for all $x \in \mathbb{R}^d$ and all $p \in \mathbb{R}^d$.

b) Assume that $f\colon \mathbb{R}^d \to \mathbb{R}$ is twice continuously differentiable and strongly convex, and assume that x^* is a minimiser of f. Show that there exists $c>0$ such that

$$f(x) \geq f(x^*)+\frac{c}{2}\|x-x^*\|^2$$

for every $x \in \mathbb{R}^d$.

Solution: An application of Taylor's theorem yields

$$f(x)=f(x^*)+\nabla f(x^*)^\top(x-x^*)+\frac{1}{2}(x-x^*)^\top \nabla^2 f(z)(x-x^*),$$

where z is on the line segment connecting x and x^*. Since x^* minimizes f, we have $\nabla f(x^*)=0$, morover an application of **a)** gives the estimate

$$(x-x^*)^\top \nabla^2 f(z)(x-x^*) \geq c\left\|x-x^*\right\|^2.$$

Therefore

$$f(x) \geq f(x^*)+\frac{c}{2}\left\|x-x^*\right\|^2.$$

[1] Not to be confused with *strictly convex*, which is a weaker notion!

c) Show that every strongly convex and continuously differentiable function $f\colon \mathbb{R}^d \to \mathbb{R}$ has a minimizer $x^* \in \mathbb{R}^d$ and that that minimizer is unique.

Solution: Let $c > 0$ be such that $g(x) = f(x) - \frac{c}{2}\|x\|^2$ is convex. Then the first order characterization of convexity implies that

$$f(x) - \frac{c}{2}\|x\|^2 = g(x) \geq g(0) + \nabla g(0)^\top x = f(0) + \nabla f(0)^\top x$$

and therefore

$$f(x) \geq f(0) + \nabla f(0)^\top x + \frac{c}{2}\|x\|^2 \geq f(0) - \|\nabla f(0)\|\,\|x\| + \frac{c}{2}\|x\|^2.$$

In particular, if $\|x\| \to \infty$, then also $f(x) \to \infty$. That is, f is coercive. Since f is also assumed to be continuosly differentiable, it is specificly continuos. It follows that f has at least one minimizer. Next assume x and y are minimizes of f. By **b)** we have that

$$f(y) \geq f(x) + \frac{c}{2}\|x - y\|^2,$$

but we also have $f(x) = f(y)$ since they are both minimizers. Thus $\|x - y\|^2 \leq 0$, which is only possible if $x = y$. This shows that the minimizer is unique.

d) Consider the *damped Newton method*

$$x_{k+1} = x_k + \alpha_k p_k \qquad \text{with} \qquad p_k = -\nabla^2 f(x_k)^{-1}\nabla f(x_k),$$

where the step length parameter α_k is chosen according to backtracking Armijo line search with parameters $\hat{\alpha} > 0$, $0 < c_1 < 1$, and $0 < \rho < 1$.

Show that p_k is a descent direction in each step, and that the sequence x_k converges to the unique minimizer x^* of f.

Hint: Use Theorem 1 in the note "Convergence of descent methods with backtracking (Armijo) linesearch…" by Anton Evgrafov.

Solution: We assume f to be a twice differentiable function. We start by showing that p_k is well defined, or more precisely, that $\nabla^2 f(x)$ is invertible for all $x \in \mathbb{R}^d$. From **a)**, we have that $p^\top \nabla^2 f(x) p \geq c\|p\|^2$ for all $p \in \mathbb{R}^d$, which implies that $\nabla^2 f(x)p$ is non-zero for all non-zero p, so $\nabla^2 f(x)$ is invertible.

That p_k is a descent direction follows from

$$\nabla f(x_k)^\top p_k = -\nabla f(x_k)^\top \nabla^2 f(x_k) \nabla f(x_k) \leq -c\|\nabla f(x_k)\|^2,$$

which implies that $\nabla f(x_k)^\top p_k < 0$. (Unless $\nabla f(x_k) = 0$, in which case x_k is the unique minimizer x^*). For the second part, define the set

$$S \coloneqq \{x \in \mathbb{R}^d | f(x) \leq f(x_0)\},$$

where x_0 is the initial value of our iteration. Note that, by the sufficient decay condition (Armijo), our iteration will never leave S. As f is continuous, S is closed. Moreover, combining the result of **b)** and **c)**, we know that there exist a minimizer $x^* \in \mathbb{R}^d$ and a constant $c > 0$ such that $f(x) \geq f(x^*) + \frac{c}{2}\|x - x^*\|^2$. Thus f is coercive and S is bounded. Finally, we introduce two constants

$$m = \inf_{x\in S} \inf_{\|p\|=1} p^\top \nabla^2 f(x) p,$$
$$M = \sup_{x\in S} \sup_{\|p\|=1} p^\top \nabla^2 f(x) p.$$

As $p^\top \nabla^2 f(x) p$ is continuous with respect to p and x, and $S \times \{p \in \mathbb{R}^d | \, \|p\| = 1\}$ is compact, we get that $M < \infty$. Furthermore, the result in **a)** implies that $m \geq c > 0$. The eigenvalues of $\nabla^2 f(x)$ are all contained in the interval $[m, M]$ whenever $x \in S$, and Theorem 1 from the note is directly applicable. We therefore know that $\|\nabla f(x_k)\| \to 0$. A quick calculation using $c > 0$ from **a)** gives

$$\begin{aligned}
\|x_k - x^*\| \, \|\nabla f(x_k)\| \geq & (x_k - x^*)^\top \nabla f(x_k) \\
= & (x_k - x^*)^\top \left(\nabla f(x_k) - \nabla f(x^*)\right) \\
= & (x_k - x^*)^\top \int_0^1 \nabla^2 f(tx_k + (1-t)x^*)(x_k - x^*) \mathrm{d}t \\
= & \int_0^1 (x_k - x^*)^\top \nabla^2 f(tx_k + (1-t)x^*)(x_k - x^*) \mathrm{d}t \\
\geq & \int_0^1 c \, \|x_k - x^*\|^2 \\
= & c \, \|x_k - x^*\|^2
\end{aligned}$$

Which implies that $\|x_k - x^*\| \leq \frac{1}{c} \|\nabla f(x_k)\|$. Combined with $\|\nabla f(x_k)\| \to 0$, this implies that $x_k \to x^*$.

3 Consider the function

$$f(x, y) = 2x^2 + y^2 - 2xy + 2x^3 + x^4.$$

a) Compute all stationary points of f and find all global or local minimisers of f.

Solution: We have

$$\nabla f = [4x - 2y + 6x^2 + 4x^3, 2y - 2x]^\top$$

and

$$\nabla^2 f = \begin{bmatrix} 4 + 12x + 12x^2 & -2 \\ -2 & 2 \end{bmatrix},$$

Hence, stationary points satisfy $y = x$ by the first component of ∇f, while the second component yields that $0 = 2x(1 + 3x + 2x^2) = x(x + 1)(2x + 1)$. Thus critical points of f are $(0, 0)$, $(-\frac{1}{2}, -\frac{1}{2})$, and $(-1, -1)$. Now,

$$\nabla^2 f(0,0) = \begin{bmatrix} 4 & -2 \\ -2 & 2 \end{bmatrix} = \nabla^2 f(-1,-1) \qquad \text{and} \qquad \nabla^2 f(-\tfrac{1}{2}, -\tfrac{1}{2}) = \begin{bmatrix} 1 & -2 \\ -2 & 2 \end{bmatrix}$$

has eigenvalues $3 \pm \sqrt{5} > 0$ and $(3 \pm \sqrt{17})/2$ (one positive, and one negative), respectively. We conclude that $(0,0)$, and $(-1,-1)$ are strict local minima, while $(-\frac{1}{2}, -\frac{1}{2})$ is a saddle point. Moreover, since $\nabla^2 f$ remains SPD both for $x > 0$ and $x < -1$ (the value of y is irrelevant), it follows that $(0,0)$ and $(-1,-1)$ are the only candidates for global minima. Evaluating $f(0,0) = 0 = f(-1,-1)$, shows that both are global minimisers of f.

b) Consider the gradient descent method with backtracking for the minimisation of f. Use the parameters $\rho = 1/2$ and $c_1 = 1/4$. Perform one step with starting value $(x_0, y_0) = (-1, 0)$. Does the method converge to a minimiser of f?

Solution: Gradient descent method gives $(x_{k+1}, y_{k+1}) = (x_k, y_k) + p_k$, with $p_k = -\nabla f_k$. Starting with preliminary step length α, $\rho = 1/2$, and $c_1 = 1/4$, we accept a new step provided

$$f\big((x_0, y_0) + \alpha p_0\big) \le f(x_0, y_0) + c\alpha \nabla f(x_0, y_0)^\top p_0 = 1 - 2\alpha$$

using that $p_0 = -\nabla f(x_0, y_0) = (2, -2)$.
Beginning with $\alpha = 1$, we reject the first try since $f\big((x_0, y_0) + \alpha p_0\big) = 13 > -1$. Reducing to $\alpha \mapsto \rho\alpha = 1/2$, still gives rejection, but $\alpha = 1/4$ succeeds, because $f\big((x_0, y_0) + \alpha p_0\big) = 1/16 \le 1/2$. Hence, we put $(x_1, y_1) = (-\frac{1}{2}, -\frac{1}{2})$,and proceed with a new round. However, (x_1, y_1) is a critical (saddle) point for f, so the gradient method stops here, thereby failing to converge to a minimiser.

c) Consider Newton's method with backtracking for the minimisation of f. Use the parameters $\rho = 1/2$ and $c_1 = 1/4$. Perform one step with starting value $(x_0, y_0) = (-1, 0)$. Does the method converge to a minimiser of f?

Solution: Similarly as in the previous exercise, the backtracking acceptance criterion for Newton's method reads

$$f\big((x_0, y_0) + \alpha p_0\big) \le f(x_0, y_0) + c\alpha \nabla f(x_0, y_0)^\top p_0 = 1 - \tfrac{1}{2}\alpha,$$

since $p_0 = -\nabla^2 f(x_0, y_0)^{-1} \nabla f(x_0, y_0) = (0, -1)$ and $c_1 = 1/4$. Starting with $\alpha = 1$, we have $f\big((x_0, y_0) + \alpha p_0\big) = 0 \le 1/2$, so the step is accepted. We then put $(x_1, y_1) = (x_0, y_0) + p_0 = (-1, -1)$. This point is a global minimiser, the conclusion being that Newton's method converged in one step.

4 (Cf. Exercise 5.1 in Nocedal & Wright)

Implement the CG method for the solution of linear systems $Ax = b$ with symmetric and positive definite matrix $A \in \mathbb{R}^{n \times n}$.

Use your method in the case where A is the Hilbert matrix, the elements of which are

$$A_{i,j} = \frac{1}{i + j - 1}, \qquad 1 \le i, j \le n.$$

Use the right hand side $b = (1, 1, \dots, 1)^T$ and the initialisation $x_0 = 0$. Test your code for dimensions $n = 5,\ 8,\ 12,\ 20$. How many iterations are required to reduce the residual below 10^{-6}? Why do your results not contradict the theoretical results concerning the CG method that were discussed in the lecture?

Hint: You might want to have a look at the condition number of the Hilbert matrix.
(Note that in MATLAB the Hilbert matrix can be produced with the command `hilb`, and in Python using `scipy.linalg.hilbert`.)

Solution: See a possible implementation on the Wiki