Norwegian University of
Science and Technology

Department of Mathematical Sciences

Examination paper for **TMA4180 Optimization I**

Academic contact during examination: Markus Grasmair

Phone:

Examination date: 14th May 2019

Examination time (from–to): 09:00–13:00

Permitted examination support material:

- Rottmann, Mathematical formulae.
- Approved basic calculator.

Other information:

- All answers should be justified and include enough details to make it clear which methods or results have been used.
- You may answer to the questions of the exam either in English or in Norwegian.
- Good luck!

Language: English

Number of pages: 10

Number of pages enclosed: 0

Checked by:

Informasjon om trykking av eksamensoppgave	
Originalen er:	
1-sidig ☐	**2-sidig** ☒
sort/hvit ☒	**farger** ☐
skal ha flervalgskjema ☐	

Date Signature

Problem 1 Consider the unconstrained optimisation problem

$$\min_{x,y} f(x,y),$$

where $f\colon \mathbb{R}^2 \to \mathbb{R}$ is defined as

$$f(x,y) = 2x^2 - 4xy + y^4 + 5y^2 - 10y.$$

a) Determine whether the function f is convex.
(5 points)

b) Find all local and global minimisers of f.
(10 points)

c) Consider now the gradient descent method and the Newton method for the solution of this optimisation problem, with step lengths chosen according to backtracking (Armijo) linesearch. Do these methods converge towards a solution of this optimisation problem? If yes, which of these two methods would you recommend, and why?
(10 points)

d) Perform one step of the gradient descent method with backtracking (Armijo) linesearch starting from the point $(x,y) = (0,0)$. Start with an initial step length $\alpha = 1$, and use the parameters $c = 0.1$ (sufficient decrease parameter) and $\rho = 0.1$ (contraction factor).
(10 points)

- *Possible solution:* We first compute

$$\nabla f(x,y) = \begin{pmatrix} 4x - 4y \\ -4x + 4y^3 + 10y - 10 \end{pmatrix}$$

and

$$\nabla^2 f(x,y) = \begin{pmatrix} 4 & -4 \\ -4 & 12y^2 + 10 \end{pmatrix}.$$

a) Since the first entry of the Hessian matrix $\nabla^2 f(x,y)$ is positive and the determinant is

$$\det \nabla^2 f(x,y) = 48y^2 + 40 - 16 = 24 + 48y^2 \geq 24 > 0,$$

it follows that f is (strongly!) convex.[1]

Alternatively, one can write

$$f(x, y) = (x - 2y)^2 + x^2 + y^2 + y^4 - 10y,$$

which is a sum of convex functions, implying that f is convex.

b) Since f is strongly convex, it follows that f has a unique local and global minimiser, which can be found by setting the gradient of f to zero. Doing this, we obtain the equations

$$\begin{aligned} 4x - 4y &= 0, \\ -4x + 4y^3 + 10y - 10 &= 0. \end{aligned}$$

From the first equation, we obtain that $x = y$. Inserting this into the second equation, we obtain the condition

$$4y^3 + 6y - 10 = 0,$$

which has the obvious solution $y = 1$. (Since f is strictly convex, this is necessarily the only solution of this equation!) Thus the unique global and local minimum is $(x, y) = (1, 1)$.

c) As shown in part a), the function f is strongly convex. Moreover, it is a polynomial and thus in particular C^2. Since f in particular is coercive, it follows that the sequence $(x_k)_{k \in \mathbb{N}}$ of iterates is bounded (and thus in particular has a convergent subsequence) and that $\nabla f(x_k) \to 0$. Since the function f only has a single critical point, this implies already that the whole sequence converges to that critical point (and global minimum). Since f is strongly convex and C^2, the same holds true for the Newton iteration.

The gradient descent method is expected to converge linearly[2] whereas the Newton method is expected to converge quadratically (as long as the line search is implemented in a reasonable way, that is, that the initial step length is $\alpha = 1$ and that $c < 0.5$). Thus it can be expected to converge in a much smaller number of steps than the gradient descent method. Moreover, each Newton step can, in this two-dimensional situation, be computed very efficiently, and thus each single Newton step is only marginally more expensive than a gradient descent step. As a consequence, I would strongly recommend the Newton method in this situation.

d) At $(x, y) = (0, 0)$ we have

$$f(0, 0) = 0 \qquad \text{and} \qquad \nabla f(0, 0) = \begin{pmatrix} 0 \\ -10 \end{pmatrix}.$$

[1] Stating that f is *strongly* (or strictly) convex is not required.

[2] Apart from very simple cases, the only situation where we can hope for a faster convergence is the case where the Hessian at the minimum is a multiple of the identity, which is (obviously) not the case in our situation.

Thus the gradient descent step is $(p, q) = \alpha(0, 10)$, and the Armijo condition (with $c = 0.1$) reads

$$f(0, 10\alpha) \leq f(0, 0) - \alpha c \|\nabla f(0, 0)\|^2 = -10\alpha.$$

For the function value at $(0, 10\alpha)$ we have the explicit expression

$$f(0, 10\alpha) = 10^4\alpha^4 + 500\alpha^2 - 100\alpha.$$

Thus the Armijo condition reads

$$10^4\alpha^4 + 500\alpha^2 \leq 90\alpha.$$

For $\alpha = 1$, this condition fails. Thus we try $\alpha = \rho \cdot 1 = 0.1$ as next possible step length. Here the left hand side becomes $10^4 \cdot 0.1^4 + 500 \cdot 0.1^2 = 6$ and the left hand side becomes $90 \cdot 0.1 = 9$. Thus the Armijo condition is satisfied and we choose the step length $\alpha = 0.1$. The next iterate in the gradient descent methods therefore becomes $(x_1, y_1) = (0, 0) + 0.1 \cdot (0, 10) = (0, 1)$.

Problem 2 We consider the constrained optimisation problem

$$-x^2 - (y - 2)^2 \to \min$$

subject to the constraint $(x, y) \in \Omega$, where the set Ω is given by the inequalities

$$y \geq 0 \qquad \text{and} \qquad x^2(x + 1) - y \geq 0.$$

a) Sketch the set Ω and determine for each point in Ω whether the LICQ holds.
(5 points)

b) Determine the tangent cone and the cone of linearised feasible directions to the set Ω in the points $(x, y) = (-1, 0)$, $(-2/3, 0)$, and $(0, 0)$.
(10 points)

c) Find all KKT points and all local and global minimisers for this optimisation problem.
(15 points)

- *Possible solution:*

  We denote in the following

$$\begin{aligned} f(x,y) &= -x^2 - (y-2)^2, \\ c_1(x,y) &= y, \\ c_2(x,y) &= x^2(x+1) - y. \end{aligned}$$

a) We first note that

$$\nabla c_1(x,y) = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad \text{and} \qquad \nabla c_2(x,y) = \begin{pmatrix} 3x^2 + 2x \\ -1 \end{pmatrix}.$$

Since none of the gradients becomes zero, it follows that the LICQ holds at every point where none or only one of the constraints is active.

Next we see that the only two points, where both constraints are active, are the points

$$(x_0, y_0) = (0,0) \qquad \text{and} \qquad (x_1, y_1) = (-1, 0).$$

At $(-1,0)$ we have

$$\nabla c_1(-1,0) = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad \text{and} \qquad \nabla c_2(-1,0) = \begin{pmatrix} 1 \\ -1 \end{pmatrix},$$

which are linearly independent; thus the LICQ holds at this point. At $(0,0)$ we have

$$\nabla c_1(0,0) = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad \text{and} \qquad \nabla c_2(0,0) = \begin{pmatrix} 0 \\ -1 \end{pmatrix},$$

which are linearly dependent; thus the LICQ fails here.

To summarise, the LICQ holds at every point in Ω apart from $(x_0, y_0) = (0,0)$.

b) At the point $(x_1, y_1) = (-1, 0)$, both constraints are active and we have

$$\begin{aligned} \mathcal{F}(-1,0) &= \Big\{(p,q) \in \mathbb{R}^2 : (p,q)\begin{pmatrix} 0 \\ 1 \end{pmatrix} \ge 0 \text{ and } (p,q)\begin{pmatrix} 1 \\ -1 \end{pmatrix} \ge 0\Big\} \\ &= \{(p,q) \in \mathbb{R}^2 : p \ge q \ge 0\}. \end{aligned}$$

Moreover, the LICQ holds, and thus $T_\Omega(-1,0) = \mathcal{F}(-1,0)$.

At the point $(x_2, y_2) = (-2/3, 0)$, only the first constraint is active. Thus

$$\mathcal{F}(-2/3, 0) = \{(p,q) \in \mathbb{R}^2 : q \ge 0\}.$$

Since, again, the LICQ holds, it follows that $T_\Omega(-2/3,0) = \mathcal{F}(-2/3,0)$.

Finally, at $(x_0, y_0) = (0,0)$ both constraints are active and we have

$$\begin{aligned}\mathcal{F}(0,0) &= \Big\{(p,q) \in \mathbb{R}^2 : (p,q)\begin{pmatrix}0\\1\end{pmatrix} \geq 0 \text{ and } (p,q)\begin{pmatrix}0\\-1\end{pmatrix} \geq 0\Big\} \\ &= \big\{(p,q) \in \mathbb{R}^2 : q = 0\big\}.\end{aligned}$$

At this point, the LICQ fails, and thus we cannot conclude that the tangent cone is equal to the cone of linearised feasible directions.

In order to compute the tangent cone, we consider first for $\tau > 0$ the sequences $(p_k, q_k) = (\pm 1/k, 0)$ and $t_k = \tau/k$. Since $c_1(p_k, q_k) = 0$ and $c_2(p_k, q_k) = (1 \pm 1/k)/k^2 \geq 0$ it follows that the sequence (p_k, q_k) is admissible and therefore

$$(\pm\tau, 0) = \lim_{k\to\infty} \frac{1}{t_k}(p_k, q_k) \in T_\Omega(0,0).$$

Moreover, the direction $(0,0)$ lies always in the tangent cone (here we can use the admissible sequence $(p_k, q_k) = (0,0)$). Thus it follows that

$$\mathcal{F}(0,0) = \big\{(p,q) \in \mathbb{R}^2 : q = 0\big\} \subset T_\Omega(0,0).$$

Since the converse inclusion $T_\Omega(0,0) \subset \mathcal{F}(0,0)$ always holds, this implies that, actually,

$$T_\Omega(0,0) = \mathcal{F}(0,0) = \big\{(p,q) \in \mathbb{R}^2 : q = 0\big\}.$$

Alternatively, one can show the converse inclusion directly by the following argument: If (p_k, q_k) is any admissible sequence approaching $(0,0)$, we have that

$$0 \leq q_k \leq p_k^2(p_k + 1) \leq 2p_k^2$$

for sufficiently large k (such that $|p_k| \leq 1$). Thus, if t_k is a sequence of positive numbers converging to zero and $(p,q) \in \mathbb{R}^d$ satisfies

$$(p,q) = \lim_{k\to\infty} \frac{1}{t_k}(p_k, q_k),$$

then

$$|q| = \lim_{k\to\infty} \frac{|q_k|}{t_k} \leq \liminf_{k\to\infty} \frac{2p_k^2}{t_k} = 2\Big(\lim_{k\to\infty} \frac{p_k}{t_k}\Big)\Big(\lim_{k\to\infty} p_k\Big) = p \cdot 0 = 0.$$

As a consequence, every tangent direction (p,q) necessarily satisfies $q = 0$ and thus $T_\Omega(0,0) \subset \{(p,q) \in \mathbb{R}^2 : q = 0\}$.

c) We note first that the points $(k, 0)$ are feasible and $f(k,0) = -k^2 - 4 \to -\infty$. Thus the problem is unbounded below and therefore does not admit a global minimum. Moreover, the point $(0,0)$ is the only point at which the LICQ fails; apart from this point, only KKT-points are candidates for local minima.

Next we compute

$$\nabla f(x,y) = \begin{pmatrix} -2x \\ 4-2y \end{pmatrix}.$$

The only point where $\nabla f(x,y) = 0$, is $(x,y) = (0,2)$, which is not feasible. Thus no point, where none of the constraints are active, can be a KKT point and thus neither a local minimum.

Next we consider the point $(x_0, y_0) = (0,0)$. Here both constraints are active and we can write

$$\nabla f(0,0) = \begin{pmatrix} 0 \\ 4 \end{pmatrix} = \lambda_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \lambda_2 \begin{pmatrix} 0 \\ -1 \end{pmatrix} = \lambda_1 \nabla c_1(0,0) + \lambda_2 \nabla c_2(0,0)$$

with Lagrange parameters $\lambda_2 \geq 0$ and $\lambda_1 = \lambda_2 + 4$. Thus $(0,0)$ is a KKT point. However, the point $(0,0)$ is no local minimiser, since the points $(1/k, 0)$, $k \in \mathbb{N}$, are feasible and

$$f(1/k, 0) = -\frac{1}{k^2} - 4 < -4 = f(0,0).$$

The next point we consider is $(x_1, y_1) = (-1, 0)$, where again both constraints are active. Here we can write

$$\nabla f(-1,0) = \begin{pmatrix} 2 \\ 4 \end{pmatrix} = 6\nabla c_1(-1,0) + 2\nabla c_2(-1,0).$$

Thus $(-1,0)$ is a KKT point with Lagrange multipliers $\lambda_1 = 6$ and $\lambda_2 = 2$. Since both Lagrange multipliers are non-zero and the constraints are linearly independent, it follows that the critical cone at $(-1,0)$ consists only of the direction $(0,0)$ and thus the second order sufficient condition is trivial satisfied. Thus the point $(-1,0)$ is a (strict and isolated) local minimum.

Now we consider the situation where only the first constraint is active, that is, $y = 0$ and $y \neq x^2(x+1)$. In this case, the KKT condition requires that

$$\nabla f(x,0) = \begin{pmatrix} -x \\ 2 \end{pmatrix} = \lambda_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \lambda_1 \nabla c_1(x,0).$$

This system can only be satisfied for $x = 0$, in which case the second constraint is also active. Thus there is no KKT point for which only the first constraint is active.

Finally, we assume that only the second constraint is active, that is $y = x^2(x+1)$ and $y > 0$. In this case, the KKT condition requires that

$$\nabla f(x,y) = \begin{pmatrix} -2x \\ 4-2y \end{pmatrix} = \lambda_2 \begin{pmatrix} 3x^2 + 2x \\ -1 \end{pmatrix}$$

for some $\lambda_2 \geq 0$. Since we require the Lagrange parameter λ_2 to be non-negative, it follows that

$$0 \leq \lambda_2 = 2y - 4$$

and thus

$$y \geq 2.$$

As a consequence it follows that

$$x^2(x+1) \geq 2,$$

which is only possible for $x \geq 1$. Since $x \neq 0$ (else both constraints are active), the equation $-2x = \lambda_2(3x^2 + 2x)$ simplifies to

$$\lambda_2 = -\frac{2}{3x+2}.$$

Since $x \geq 1$, this implies that $\lambda_2 < 0$. Thus there is no KKT point with only the second constraint active.

To summarise, there are two KKT points, $(x_0, y_0) = (0,0)$ and $(x_1, y_1) = (-1, 0)$, the second of which is the only local, but not global, solution of the constrained optimisation problem.

Problem 3 For a fixed $t \in \mathbb{R}$, we consider the *elastic net* optimisation problem

$$f(x) = \frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |x| \to \min. \qquad (P)$$

An equivalent formulation (in the sense that x^* solves (P) if and only if $x^* = y^*$ solve (P')) is the problem

$$\min_{x,y} \frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |y| \to \min \qquad \text{subject to } x = y. \qquad (P')$$

a) Formulate the Lagrangian dual of the problem (P') as a constrained optimisation problem.
(10 points)

b) Find an explicit formula for the solution of (P) depending on t.
(10 points)

- *Possible solution:*

a) The Lagrangian for the problem (P') is

$$\mathcal{L}(x,y;\lambda) = \frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |y| - \lambda(x-y).$$

Thus the dual problem is

$$\max_{\lambda\in\mathbb{R}} \min_{x,y} \Big(\frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |y| - \lambda(x-y)\Big).$$

In order to find a reasonable formulation for this problem, we compute the (extended real-valued) function $q\colon \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$,

$$\begin{aligned} q(\lambda) &= \min_{x,y} \Big(\frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |y| - \lambda(x-y)\Big) \\ &= \min_{x} \Big(\frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 - \lambda x\Big) + \min_{y}(|y| + \lambda y). \end{aligned}$$

The minimum for the first term is attained for $x - t + x - \lambda = 0$, that is,

$$x = (\lambda + t)/2,$$

at which point the value is

$$\begin{aligned} \min_{x} \Big(\frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 - \lambda x\Big) &= \frac{1}{2}\Big(\frac{\lambda - t}{2}\Big)^2 + \frac{1}{2}\Big(\frac{\lambda + t}{2}\Big)^2 - \lambda\frac{\lambda + t}{2} \\ &= \frac{t^2}{4} - \frac{\lambda^2}{4} - \frac{\lambda t}{2} \\ &= -\frac{1}{4}(\lambda + t)^2 + \frac{t^2}{2}. \end{aligned}$$

For the second term we obtain

$$\min_{y}(|y| + \lambda y) = \begin{cases} 0 & \text{if } |\lambda| \le 1, \\ -\infty & \text{if } |\lambda| > 1, \end{cases}$$

Thus

$$q(\lambda) = \begin{cases} -\frac{1}{4}(\lambda + t)^2 + \frac{t^2}{2} & \text{if } |\lambda| \le 1, \\ -\infty & \text{if } |\lambda| > 1. \end{cases}$$

Thus we can write the dual problem as

$$\max_{\lambda} \Big(-\frac{1}{4}(\lambda + t)^2 + \frac{t^2}{2}\Big) \qquad \text{subject to } |\lambda| \le 1. \qquad (D)$$

b) Since the function

$$g(x,y) = \frac{1}{2}(x-t)^2 + \frac{1}{2}x^2 + |y|$$

is convex and the constraint $x = y$ is linear, strong duality holds. In particular, this implies that, if (x^*, y^*) solve (P') and λ^* solves the dual problem (D), then (x^*, y^*) solve

$$\min_{x,y} \mathcal{L}(x,y;\lambda^*). \tag{1}$$

Moreover, as has already been computed in the first part of this problem, the x-coordinate of the solution of $\min_{x,y} \mathcal{L}(x,y;\lambda^*)$ is

$$x^* = \frac{\lambda^* + t}{2}.$$

The dual problem (D) is (apart from constants and constant factors) simply the projection of $-t$ to the interval $[-1,1]$. Thus the solution is

$$\lambda^* = \max\{-1, \min\{1, -t\}\}.$$

As a consequence, we obtain

$$x^* = \frac{\lambda^* + t}{2} = \frac{\max\{-1, \min\{1,-t\}\} + t}{2} = \begin{cases} \dfrac{t-1}{2} & \text{if } t \geq 1, \\ 0 & \text{if } -1 \leq t \leq 1, \\ \dfrac{t+1}{2} & \text{if } t \leq -1. \end{cases}$$

Problem 4 We consider the optimisation problem

$$f(x) \to \min \qquad \text{s.t. } c(x) \geq 0,$$

where $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex and $\mathcal{C}^1$, and $c\colon \mathbb{R}^d \to \mathbb{R}$ is concave and $\mathcal{C}^1$. Moreover, we assume that the feasible set $\Omega = \left\{x \in \mathbb{R}^d : c(x) \geq 0\right\}$ is non-empty.

a) Formulate the KKT conditions for this optimisation problem and state (with a brief explanation) if they are sufficient and/or necessary optimality conditions.
(5 points)

b) Show that in this situation Slater's constraint qualification is satisfied, if and only if every point $x \in \Omega$ satisfies the LICQ.
(10 points)

- *Possible solution:*

a) The KKT conditions read

$$\begin{aligned} \nabla f(x) &= \lambda \nabla c(x), \\ c(x) &\geq 0, \\ \lambda &\geq 0, \\ \lambda &= 0 \qquad \text{if } c(x) = 0. \end{aligned}$$

Since we have a convex optimisation problem with a concave constraint, it follows that the KKT conditions are sufficient optimality conditions. However, they are not necessary, unless in addition some constraint qualification is satisfied.

(Consider for instance the (univariate) problem $f(x) = x \to \min$ subject to $c(x) = -x^2 \geq 0$. Here the points $x = 0$ is the only feasible point and thus the unique global solution of the problem, but $f'(0) = 1$ cannot be written as a positive multiple of $c'(0) = 0$.)

b) We recall first that Slater's constraint qualification requires the existence of $x \in \mathbb{R}^d$ with $c(x) > 0$, whereas the LICQ requires that at each point x with $c(x) = 0$ we have $\nabla c(x) \neq 0$ (the gradient being zero at a point where the constraint c is active is the only possibility for the LICQ to fail).

Assume now that Slater's constraint qualification holds. That is, there exists $x_0 \in \mathbb{R}^d$ with $c(x_0) > 0$. Then, if $x \in \mathbb{R}^d$ satisfies $c(x) = 0$, then x is no maximum of c. Since c is concave (or: $-c$ is convex), this implies that $\nabla c(x) \neq 0$, and thus the LICQ holds at x.

Conversely, assume that Slater's constraint qualification does not hold. Then there exists no $x \in \mathbb{R}^d$ with $c(x) > 0$, which implies that every point x with $c(x) = 0$ is a maximum of c (such points exist, as Ω is non-empty). As a consequence we have $\nabla c(x) = 0$ at every such point (the first order optimality condition holds), and thus the LICQ fails.