Problem 1 Consider the function $f\colon \mathbb{R}^2 \to \mathbb{R}$ given by

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

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

We start with computing the stationary points of f, which are the solutions of the equation $\nabla f(x,y) = 0$. First,

$$\nabla f(x,y) = \begin{pmatrix} 4x^3 - 6x^2 + 4x - 2y \\ -2x + 2y \end{pmatrix}.$$

Since the second line of this expression has to be zero, we immediately obtain the condition $x = y$. Inserting this into the first line and setting this line to zero, we obtain the equation

$$4x^3 - 6x^2 + 2x = 0.$$

Obviously, one of the solutions of this equation is $x = 0$. The other two solutions, $x = 1$ and $x = 1/2$, can then be obtained either by solving the ensuing quadratic polynomial or by guessing (recall that all rational roots of polynomials with rational coefficients can be "guessed" methodically). We therefore obtain the three stationary points

$$\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad \begin{pmatrix} x_3 \\ y_3 \end{pmatrix} = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}.$$

Next we determine, which of the stationary points are local minima. To that end we compute the Hessian of f:

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

Specifically, we have

$$\nabla^2 f(0,0) = \nabla^2 f(1,1) = \begin{pmatrix} 4 & -2 \\ -2 & 2 \end{pmatrix}.$$

The upper left entry of this matrix is 4, and its determinant is also 4. Since both of these values are positive, it follows that the matrix is positive definite. This implies that the points $(0,0)$ and $(1,1)$ are local minima of f.

Next we see that

$$\nabla^2 f(1/2,1/2) = \begin{pmatrix} 1 & -2 \\ -2 & 2 \end{pmatrix}.$$

The determinant of this matrix is negative, and therefore it is not positive semi-definite, which in turn implies that the point $(1/2,1/2)$ is no local minimum.[1]

Therefore the only local minima of f are $(0,0)$ and $(1,1)$.

[1] If one would analyse this matrix further, one would see that it is actually indefininte or, more precisely, it has both a positive and a negative eigenvalue. As a consequence, the point $(1/2,1/2)$ is a saddle point of f. (This analysis was, however, not part of the problem.)

Finally, we study the question of global minimizers. To that end we note first that f is continuous (which is obvious) and coercive. The latter can be seen, for instance, by writing

$$f(x, y) = x^2(x-1)^2 + (x-y)^2. \tag{1}$$

The first term of this expression tends to ∞ as $x \to \pm\infty$. Since the second term is non-negative, this implies that $f(x, y) \to \infty$ whenever $x \to \pm\infty$. Moreover, if $y \to \infty$ and x stays bounded, the second term in (1) tends to ∞, while the first term remains bounded. This argumentation shows that $f(x, y) \to \infty$ whenever at least one of the arguments tends to ∞. As a consequence, f is coercive. Since f is continuous and coercive, we know that at least one global minimizer exists. The only candidates for a global minimizer are the two local minimizers, that is, the points $(0, 0)$ and $(1, 1)$. Since $f(0, 0) = f(1, 1)$, it follows that, actually, both of these points are global minimizers. Thus the global minimizers of f are $(0, 0)$ and $(1, 1)$.[2]

b) Determine whether the function f is convex or not.

There are several different possibilities for showing that the function f is non-convex, for instance the following four:

- The Hessian of f at $(1/2, 1/2)$ (which we have computed above) is not positive semi-definite.
- The function f has a stationary point that is not a global minimizer (the point $(1/2, 1/2)$).
- The function f has exactly two global minimizers. Since the set of global minimizers of a convex function is convex, it follows that f is non-convex.
- We have

$$f(1/2, 1/2) = \frac{1}{16} > \frac{1}{2} f(0, 0) + \frac{1}{2} f(1, 1) = 0.$$

c) Starting at the point $(x, y) = (0, 1)$ compute one step of the steepest descent method with backtracking (Armijo) linesearch (see Algorithm 3.1 in Nocedal and Wright). Start with an initial step length $\bar{\alpha} = 1$ and use the parameters $c = 0.25$ (sufficient decrease parameter) and $\rho = 0.1$ (contraction factor).

At $(x, y) = (0, 1)$ we have

$$\nabla f(0, 1) = \begin{pmatrix} -2 \\ 2 \end{pmatrix} \qquad \text{and therefore} \qquad p = \begin{pmatrix} 2 \\ -2 \end{pmatrix}.$$

[2] Another, simpler, approach to this problem is to note that f is non-negative (which can be seen from (1)) and that $f(x, y) = 0$ if and only if either $(x, y) = (0, 0)$ or $(x, y) = (1, 1)$. This shows directly that $f(0, 0) = f(1, 1) \le f(x, y)$ for all $(x, y) \notin \{(0, 0), (1, 1)\}$, proving that $(0, 0)$ and $(1, 1)$ are global minimizers of f.

The Armijo condition now requires that

$$f((0,1)+\alpha p) \le f(0,1) + \alpha\frac{1}{4}\nabla f(0,1)^T p = 1 - \alpha\frac{1}{4}\cdot 8 = 1 - 2\alpha.$$

With $\alpha = 1$ we obtain

$$f((0,1)+\alpha p) = f(2,-1) = 13 \not\le 1 - 2\alpha = -1.$$

Thus we try the step size $\alpha = \rho \cdot 1 = 0.1$ and obtain

$$f((0,1)+\alpha p) = f(0.2, 0.8) = 0.3856 \le 1 - 2\alpha = 0.8.$$

Thus the step size $\alpha = 0.1$ is accepted and the next point in the iteration would be $(x_1, y_1) = (0.2, 0.8)$.

Problem 2 Find all parameters $\alpha \in \mathbb{R}$ for which the point $(x,y) = (3,1)$ is a local solution of the optimization problem

$$x + \alpha y \to \min$$

subject to the constraints

$$\begin{aligned} xy - 3 &\ge 0,\\ 10 - x^2 - y^2 &\ge 0. \end{aligned}$$

Denote $c_1(x,y) = xy - 3$ and $c_2(x,y) = 10 - x^2 - y^2$ the two constraints in the optimization problem in question. At the point $(x,y) = (3,1)$ both of these constraints are active. Moreover,

$$\nabla c_1(3,1) = \begin{pmatrix}1\\3\end{pmatrix} \qquad \text{and} \qquad \nabla c_2(3,1) = -\begin{pmatrix}6\\2\end{pmatrix},$$

which are obviously linearly independent vectors. Thus the LICQ holds at $(3,1)$ and therefore the KKT conditions are necessary for $(3,1)$ to be a local minimizer of this problem. Ignoring the conditions $c_1(3,1) \ge 0$ and $c_2(3,1) \ge 0$ (which we have already shown to hold), these conditions read as

$$\begin{aligned} \begin{pmatrix}1\\ \alpha\end{pmatrix} - \lambda_1\begin{pmatrix}1\\3\end{pmatrix} + \lambda_2\begin{pmatrix}6\\2\end{pmatrix} &= 0,\\ \lambda_1 &\ge 0,\\ \lambda_2 &\ge 0. \end{aligned}$$

The first line in the equation above implies that

$$\lambda_1 = 1 + 6\lambda_2.$$

Inserting this into the second line in the equation, we obtain that

$$\alpha = 3\lambda_1 - 2\lambda_2 = 16\lambda_2 + 3.$$

This implies that $\lambda_2 \geq 0$ if and only if $\alpha \geq 3$. Since $\lambda_1 = 1 + 6\lambda_2$, it follows that, in this case, also $\lambda_1 \geq 0$. Thus the KKT conditions are satisfied if and only if $\alpha \geq 3$. Additionally, the Lagrange parameters are both strictly positive if $\alpha > 3$.

In order to find out whether the point $(3, 1)$ is actually a local minimizer, we test the second order sufficient conditions.[3] For $\alpha > 3$, the second order sufficient conditions are trivially satisfied, because in this case all the Lagrange parameters are positive, and therefore the critical cone contains only the zero-vector.

For $\alpha = 3$ we have the Lagrange parameters $\lambda_1 = 1$ and $\lambda_2 = 0$. Thus, a vector $w \in \mathbb{R}^2$ is contained in the critical cone $\mathcal{C}(3, 1; 1, 0)$, if and only if

$$w^T \begin{pmatrix} 1 \\ 3 \end{pmatrix} = 0 \qquad \text{and} \qquad w^T \begin{pmatrix} -6 \\ -2 \end{pmatrix} \geq 0$$

(the vectors appearing in these conditions are simply the gradients of the constraints at $(x, y) = (3, 1)$; the inequality in the second condition is due to the fact that $\lambda_2 = 0$). The first condition implies that w is of the form $w = \mu(-3, 1)$ for some $\mu \in \mathbb{R}$. The second condition then implies that

$$0 \leq \mu \begin{pmatrix} -3 & 1 \end{pmatrix} \begin{pmatrix} -6 \\ -2 \end{pmatrix} = 16\mu,$$

that is, $\mu \geq 0$. Thus

$$\mathcal{C}(3, 1; 1, 0) = \{\mu(-3, 1) : \mu \geq 0\}.$$

Next we compute

$$\nabla^2 \mathcal{L}(3, 1; 1, 0) = -\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Thus, if $w \in \mathcal{C}(3, 1; 1, 0) \setminus \{0\}$, that is, if $w = \mu(-3, 1)$ with $\mu > 0$, we obtain that

$$w^T \nabla^2 \mathcal{L}(3, 1; 1, 0) w = -\mu^2 \begin{pmatrix} -3 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} -3 \\ 1 \end{pmatrix} = -\mu^2 \begin{pmatrix} -3 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -3 \end{pmatrix} = 6\mu^2 > 0.$$

Thus the second order sufficient condition is satisfied also for $\alpha = 3$, and therefore $(3, 1)$ is a local minimizer of the constrained optimization problem in question, if and only if $\alpha \geq 3$.

[3] Alternatively, it would also be possible to note that we are minimizing a convex functional and that the constrained set is locally convex around the point $(3, 1)$ (the constrained set consists of two convex regions, one in the quadrant $x > 0$ and $y > 0$, the other one in the quadrant $x < 0$ and $y < 0$). In such a (basically convex) case, the first order necessary condition is actually sufficient for a local minimizer. However, we did not prove this fact in the lecture, and it appears not to be mentioned in Nocedal & Wright as well.

Problem 3 Consider the optimization problem

$$x^2 + 2xy + 2y^2 \to \min \qquad \text{subject to} \qquad x + y - 1 = 0. \tag{2}$$

The unique (local and global) solution of this problem is the point $(x, y) = (1, 0)$ (you don't have to show this).

a) Formulate the unconstrained optimization problem that results from the application of the quadratic penalty method to (2), and compute the solution for all possible penalty parameters $\mu > 0$.

The quadratic penalty method requires the solution of the unconstrained optimization problem

$$\mathcal{Q}(x, y; \mu) = x^2 + 2xy + 2y^2 + \frac{\mu}{2}(x + y - 1)^2 \to \min .$$

It is easy to see that the function $\mathcal{Q}$ is convex for every $\mu > 0$, as it is the sum of convex functions. Thus the necessary and sufficient condition for a global minimizer of this function is the equation

$$\nabla\mathcal{Q}(x, y; \mu) = \begin{pmatrix} 2x + 2y + \mu(x + y - 1) \\ 2x + 4y + \mu(x + y - 1) \end{pmatrix} = 0.$$

Subtracting the first line from the second line, we immediately see that $y = 0$. Then the first line simplifies to

$$(2 + \mu)x - \mu = 0 \qquad \text{or} \qquad x = \frac{\mu}{2 + \mu}.$$

Thus the solution is $(x, y) = (\frac{\mu}{2+\mu}, 0)$.

b) Formulate the augmented Lagrangian $\mathcal{L}_A(x, \lambda; \mu)$ corresponding to (2) and compute its minimizers for all possible parameters $\lambda \in \mathbb{R}$ and $\mu > 0$. For which parameters does the minimizer of the augmented Lagrangian coincide with the solution of (2)?

The augmented Lagrangian reads as

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

Again, this function is convex, and the necessary and sufficient condition for a minimizer is the equation

$$\nabla\mathcal{L}_A(x, y, \lambda; \mu) = \begin{pmatrix} 2x + 2y - \lambda + \mu(x + y - 1) \\ 2x + 4y - \lambda + \mu(x + y - 1) \end{pmatrix} = 0.$$

Again we see immediately that $y = 0$ and the first line simplifies to

$$(2 + \mu)x - \lambda - \mu = 0 \qquad \text{or} \qquad x = \frac{\lambda + \mu}{2 + \mu}.$$

The solution is therefore $(x, y) = (\frac{\lambda+\mu}{2+\mu}, 0)$.
Moreover, we see that the solution is exact (that is, coincides with the solution of the original problem), if and only if $\frac{\lambda+\mu}{2+\mu} = 1$, which is the case if and only if $\lambda = 2$.[4]

Problem 4 The set $\Omega \subset \mathbb{R}^2$ is given by the constraints

$$\begin{aligned} x + 1 &\geq 0, \\ 1 - x - y &\geq 0, \\ (x+1)^2 y^3 &\geq 0. \end{aligned}$$

Using a set of suitable linear inequalities and equalities, describe both the tangent cone and the cone of linearized feasible directions for Ω at $(x, y) = (1, 0)$.

The set Ω is depicted here:

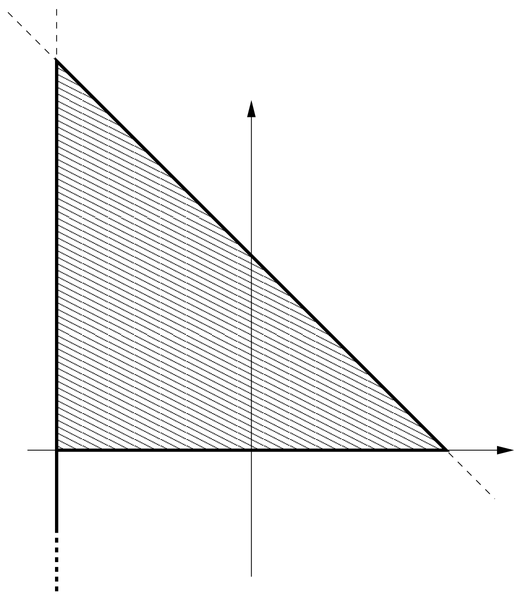

At the point $(x, y) = (1, 0)$ only the second and the third condition are active, which means that we can (or rather: have to) ignore the first condition for the rest of this example.

We start with computing the cone of linearized feasible directions. To that end we need to compute the gradients of the (relevant) constraint functions $c_2(x, y) = 1 - x - y$ and $c_3(x, y) = (x+1)^2 y^3$ at $(x, y) = (1, 0)$. We have

$$\nabla c_2(x, y) = -\begin{pmatrix} 1 \\ 1 \end{pmatrix} \qquad \text{and} \qquad \nabla c_3(x, y) = \begin{pmatrix} 2(x+1)y^3 \\ 3(x+1)^2 y^2 \end{pmatrix}.$$

[4] Note that $\lambda = 2$ is also the Lagrange multiplier corresponding to the solution of the constrained optimization problem.

At the point $(1,0)$ this reduces to

$$\nabla c_2(1,0) = -\begin{pmatrix}1\\1\end{pmatrix} \qquad \text{and} \qquad \nabla c_3(1,0) = \begin{pmatrix}0\\0\end{pmatrix}.$$

The cone of linearized feasible directions is defined as

$$\mathcal{F}(1,0) = \{(u,v) \in \mathbb{R}^2 : (u,v)\nabla c_k(1,0) \geq 0 \text{ for } k = 2,3\}.$$

Since $\nabla c_3(1,0) = 0$, the second inequality is trivially satisfied. Thus the cone of linearized feasible directions is characterized by the single linear inequality

$$\begin{pmatrix}u & v\end{pmatrix} - \begin{pmatrix}1\\1\end{pmatrix} = -u - v \geq 0.$$

Since the LICQ is (obviously) not satisfied at $(1,0)$, the tangent cone may be (and actually is) different from $\mathcal{F}(1,0)$. In order to compute the tangent cone, we note that, locally around the point $(x,y) = (1,0)$, the condition $(x+1)^2y^3 \geq 0$ is equivalent to $y \geq 0$. Thus the constrained set can, in a neighborhood of $(1,0)$ be equivalently described by the inequalities

$$1 - x - y \geq 0 \qquad \text{and} \qquad y \geq 0.$$

Note that this rewriting of the constraints does not change the tangent cone at the point $(1,0)$. The cone of linearized feasible directions for this set of constraints is given by the inequalities

$$\begin{pmatrix}u & v\end{pmatrix}\begin{pmatrix}-1\\-1\end{pmatrix} \geq 0 \qquad \text{and} \qquad \begin{pmatrix}u & v\end{pmatrix}\begin{pmatrix}0\\1\end{pmatrix} \geq 0.$$

Since the vectors $(-1,-1)$ and $(0,1)$ are linearly independent, the LICQ holds and thus this cone of linearized feasible directions coincides with the tangent cone. Therefore it follows that the tangent cone to Ω consists of all $(u,v) \in \mathbb{R}^2$ satisfying the inequalities

$$-u - v \geq 0 \qquad \text{and} \qquad v \geq 0.$$

It is also possible (but much more annoying) to determine the tangent cone by only using its definition. Here we note first that $T_\Omega(1,0) \subset \mathcal{F}(1,0)$, which implies that the condition $-u-v \geq 0$ is a necessary condition for a vector (u,v) to be contained in $T_\Omega(1,0)$. Next we show that the inequality $v \geq 0$ is necessary. To that end we assume that $(u,v) \in T_\Omega(1,0)$. Then there exists sequences $(u_k, v_k) \in \Omega$ converging to $(1,0)$ and $t_k \in \mathbb{R}_{>0}$ such that $(u_k - 1, v_k)/t_k \to (u,v)$. Since $(u_k, v_k) \to (1,0)$ it follows from the definition of Ω that $v_k \geq 0$ for k sufficiently large. Thus $v = \lim_{k\to\infty} v_k/t_k \geq 0$.

We have thus shown that the inequalities $-u - v \geq 0$ and $v \geq 0$ are necessary for (u,v) to be contained in $T_\Omega(1,0)$. It remains to show that they are also sufficient. To that end assume that $(u,v) \in \mathbb{R}^2$ satisfies $u \leq v$ and $v \geq 0$. Define $u_k = 1 + u/k$, $v_k = v/k$, and $t = 1/k$. Then $(u_k, v_k) \to (1,0)$ and $(u_k, v_k) \in \Omega$ for k sufficiently large (more precisely: for $k \geq |u|/2$). Since $((u_k, v_k) - (1,0))/t_k = (u,v)$ for all k, this implies that $(u,v) \in T_\Omega(1,0)$, showing the sufficiency of these inequalities.

Problem 5 Show that a (not necessarily differentiable) function $f\colon \mathbb{R}^n \to \mathbb{R}_{>0}$ is convex, if the function $x \mapsto \ln(f(x))$ is convex.

Define $g\colon \mathbb{R}^n \to \mathbb{R}$, $g(x) = \ln(f(x))$. Then $f(x) = \exp(g(x))$.

Assume that x, $y \in \mathbb{R}^n$ and $0 \le \lambda \le 1$. Since g is convex it follows that

$$g(\lambda x + (1-\lambda)y) \le \lambda g(x) + (1-\lambda)g(y).$$

Now the fact that exp is monotonically increasing implies that

$$f(\lambda x + (1-\lambda)y) = \exp(g(\lambda x + (1-\lambda)y)) \le \exp(\lambda g(x) + (1-\lambda)g(y)).$$

Finally we use the convexity of exp to obtain that

$$\begin{aligned} f(\lambda x + (1-\lambda)y) &\le \exp(\lambda g(x) + (1-\lambda)g(y)) \\ &\le \lambda \exp(g(x)) + (1-\lambda)\exp(g(y)) = \lambda f(x) + (1-\lambda)f(y). \end{aligned}$$

Problem 6 Assume that the sequence $(x_k)_{k\in\mathbb{N}}$ is generated by the gradient descent method with backtracking linesearch for the minimization of a function f, and that $\nabla f(x_k) \neq 0$ for all k. Assume moreover that $\bar{x}$ is an accumulation point of the sequence $(x_k)_{k\in\mathbb{N}}$. Show that $\bar{x}$ is not a local maximum of f.

Since the sequence x_k is generated using a backtracking linesearch method, it satisfies the Armijo condition

$$f(x_{k+1}) = f(x_k + \alpha_k p_k) \le f(x_k) + c\alpha_k \nabla f(x_k)^T p_k.$$

With $p_k = -\nabla f(x_k) \neq 0$ implies that

$$f(x_{k+1}) \le f(x_k) - c\alpha_k \|\nabla f(x_k)\|^2 < f(x_k).$$

In other words, the sequence $f(x_k)$ is strictly decreasing. Now assume that $\bar{x}$ is an accumulation point of the sequence x_k. Then there exists a subsequence $x_{k'}$ converging to $\bar{x}$. Since $f(x_k)$ is strictly decreasing, it follows that also $f(x_{k'})$ is strictly decreasing. Since $f(x_{k'}) \to f(\bar{x})$, this implies that $f(x_{k'}) > f(\bar{x})$ for every k'. Thus we have shown that there exists a sequence $x_{k'}$ converging to $\bar{x}$ such that $f(x_{k'}) > f(\bar{x})$ for every k', which in turn shows that $\bar{x}$ is no local maximum of f.

Problem 7 We consider a line search method of the form $x_{k+1} = x_k + \alpha_k p_k$ for the minimization of the function $f\colon \mathbb{R}^n \to \mathbb{R}$, where the search direction p_k is given as

$$p_k = -\operatorname{sgn}((\nabla f(x_k))_i)\, e_i,$$

where the index i is chosen such that $|(\nabla f(x_k))_i|$ is maximal. Here e_i with $1 \le i \le n$ denotes the i-th standard basis vector in $\mathbb{R}^n$.

a) Show that the direction p_k is a descent direction whenever x_k is no stationary point of f.

We recall that p_k is a descent direction for f at x_k, if and only if $p_k^T \nabla f(x_k) < 0$. Assume now that x_k is no stationary point of f, that is, $\nabla f(x_k) \neq 0$. Since the index i in the definition of p_k is chosen in such a way that $|(\nabla f(x_k))_i|$ is maximal, we obtain in particular that $|(\nabla f(x_k))_i| > 0$. Thus

$$\begin{aligned} p_k^T \nabla f(x_k) &= -\operatorname{sgn}((\nabla f(x_k))_i)\, e_i^T \nabla f(x_k) \\ &= -\operatorname{sgn}((\nabla f(x_k))_i)(\nabla f(x_k))_i = -|(\nabla f(x_k))_i| < 0, \end{aligned}$$

which shows that p_k is a descent direction.

b) Assume that f is twice continuously differentiable and coercive and that the step lengths α_k satisfy the Wolfe conditions (see Nocedal & Wright, equation (3.6)). Show that $\lim_{k\to\infty} \|\nabla f(x_k)\| = 0$.

We apply Zoutendijk's result (Theorem 3.2 in Nocedal & Wright) to obtain that

$$\sum_{k=0}^{\infty} \cos^2 \vartheta_k \|\nabla f(x_k)\|^2 < \infty$$

with

$$\cos^2 \vartheta_k = \frac{(p_k^T \nabla f(x_k))^2}{\|p_k\|^2 \|\nabla f(x_k)\|^2}.$$

Inserting the definition of $\cos^2 \vartheta_k$ in the infinite sum above and noting that $\|p_k\| = 1$ for every k, we obtain that

$$\sum_{k=0}^{\infty} (p_k^T \nabla f(x_k))^2 < \infty.$$

Now note that

$$|p_k^T \nabla f(x_k)| = \max_{1 \le i \le n} |(\nabla f(x_k))_i|.$$

Thus

$$\sum_{k=0}^{\infty} \big(\max_{1 \le i \le n} |(\nabla f(x_k))_i|\big)^2 < \infty.$$

This in particular implies that

$$\max_{1\leq i\leq n} |(\nabla f(x_k))_i| \to 0$$

(for a series to converge it is necessary that its terms converge to zero). Now the estimate

$$\|\nabla f(x_k)\| \leq \sqrt{n} \max_{1\leq i\leq n} |(\nabla f(x_k))_i|$$

implies the convergence of $\|\nabla f(x_k)\|$ to zero.