THEORY OF UNCONSTRAINED OPTIMISATION

PRELIMINARY VERSION

MARKUS GRASMAIR

In this note, we will discuss the basic theory required for the solution of finite dimensional, continuous optimisation problems. In the first section, we will discuss the basic definitions and show conditions that guarantee existence of solutions to optimisation problems, both in the constrained and unconstrained case. Then we restrict ourselves to unconstrained problems, where we derive different necessary and sufficient optimality conditions. Here, the constrained case is significantly more difficult and will be discussed later in the lecture. Finally, we consider in particular convex functions and their properties as relates to optimisation. It turns out that convex functions are in a sense among the most simple functions to minimise, both in a theoretical setting (as discussed here) and in practice with numerical algorithms (as we will show throughout the course).

1. Solutions of Optimisation Problems

We will start with discussing the existence of solutions of minimisation problems of the form

$$\min_{x \in \Omega} f(x),$$

where $f : \mathbb{R}^d \to \mathbb{R}$ is a real valued function (the *cost function* or *objective function*) and $\Omega \subset \mathbb{R}^d$ is some set (the *feasible set*).

1.1. Notions of Minimisers.

First we have to clarify what we mean by a solution of an optimisation problem.

Definition 1.1 (Global minimiser)**.** A point $x^* \in \Omega$ is called a *global minimiser* (or *global minimum*, or *global solution*) of the optimisation problem $\min_{x \in \Omega} f(x)$, if

$$f(x^*) \le f(x)$$

for all $x \in \Omega$.

The point x^* is *strict global minimiser*, if $f(x^*) < f(x)$ for all $x \in \Omega$, $x \neq x^*$.

Example 1.2. The point $x = -\pi/2$ is a global minimiser of the function $f(x) = \sin(x)$, since $\sin(-\pi/2) = -1$ and $\sin(x) \ge -1$ for every $x \in \mathbb{R}$. However, we also have that $\sin(3\pi/2) = -1 = \sin(-\pi/2)$. Thus $x = -\pi/2$ is a global minimiser, but not a strict global minimiser of f.

One problem of global minimisers is that they are incredibly hard to recognise in general. In order to verify that a point x^* is a global minimiser, one would have to compare $f(x^*)$ with every other value $f(x)$, no matter how large the distance between x and x^* is. In actual applications, however, one usually may only obtain the value of f (and, possibly, some of its derivatives) at a small number of selected points. With only this information available, only in very special cases is it possible to prove that a given point x^* is really a global minimiser.

As an alternative, we therefore consider local minimisers:

Date: 22nd February 2023.

Definition 1.3 (Local minimiser)**.** A point $x^* \in \Omega$ is called a *local minimiser* (or *local minimum*, or *local solution*) of the optimisation problem $\min_{x\in\Omega} f(x)$, if there exists $\varepsilon > 0$ such that $f(x^*) \leq f(x)$ whenever $x \in \Omega$ satisfies $\|x - x^*\| \leq \varepsilon$.

Slightly strengthening this notation, we obtain:

Definition 1.4 (Strict local minimiser)**.** A point x^* is called a *strict local minimiser* of $\min_{x\in\Omega} f(x)$, if there exists $\varepsilon > 0$ such that

$$f(x^*) < f(x)$$

whenever $x \in \Omega$, $x \neq x^*$ satisfies $\|x - x^*\| \leq \varepsilon$.

That is, we replace the inequality $\leq$ by the strict inequality $<$ in the definition of the local minimiser.

In addition, it makes sometimes sense to strengthen this notion further:

Definition 1.5 (Isolated local minimiser)**.** A point $x^* \in \Omega$ is called an *isolated local minimiser* of the problem $\min_{x\in\Omega} f(x)$, if there exists $\varepsilon > 0$ such that x^* is the only local minimiser of f in Ω an ε-ball around x^*. That is, if $y^* \neq x^*$ is another local minimiser of f in Ω, then $\|x^* - y^*\| > \varepsilon$.

Example 1.6. The point $x = -\pi/2$ is a local minimiser of the function $f(x) = \sin(x)$. In fact, it is an isolated local minimiser, because it is the only local minimiser of f within the open interval $(-5\pi/2, 3\pi/2)$.

Example 1.7. Define $f\colon \mathbb{R}^2 \to \mathbb{R}$, $f(x,y) = (x^2 - y)^2$. Then $f(x,y) \geq 0$ for all $(x,y) \in \mathbb{R}^2$, and $f(x,y) = 0$ if $x^2 = y$. Thus all points of the form (x, x^2) for $x \in \mathbb{R}$ are local and global minimisers of f, but none of them is an isolated minimiser.

Example 1.8. From the definition, it is easy to see that every isolated local minimiser is necessarily a strict local minimiser. The converse, however, does not necessarily hold as seen by the (rather pathological) function

$$f(x) = \begin{cases} 2x^2 + x^2\sin(1/x) & \text{if } x \neq 0, \\ 0 & \text{if } x = 0. \end{cases}$$

This function has a strict local minimiser at $x = 0$ (which is at the same time the unique global minimiser of f), but there exists a sequence of (isolated!) local minimisers converging to 0. Thus the minimiser at 0 is not isolated. See also Figure 1.

1.2. **Existence of minimisers.** In general, an optimisation problem need not attain local or global solutions. Consider for instance the following optimisation problems (see also Figure 2):

- Minimise the function $f(x) = 1/x$ for $x \in \mathbb{R} \setminus \{0\}$, $f(0) = 1$.
  Here the function f is unbounded below (we have $\lim_{x\to 0^-} f(x) = -\infty$), and thus it cannot have a global minimiser.
- Minimise the function $f(x) = e^{-x^2}$ for $x \in \mathbb{R}$.
  The function f satisfies $f(x) > 0$ for all $x \in \mathbb{R}$, but $\lim_{x\to\pm\infty} f(x) = 0$. That is, the function values can be arbitrarily close to 0, but there is no point $x \in \mathbb{R}$ at which the value 0 is actually attained.
- Minimise the function $f(x) = x$ for $x > 0$ and $f(x) = x^2 + 1$ for $x \leq 0$.
  Again, the function attains only positive values that can be arbitrarily close to 0.

We have seen above that an optimisation problem need not necessarily have a solution: As seen above, the function $f(x) = e^{-x^2}$ does not attain a minimum, and nor does the function f defined by $f(x) = x$ if $x > 0$ and $f(x) = x^2 + 1$ if

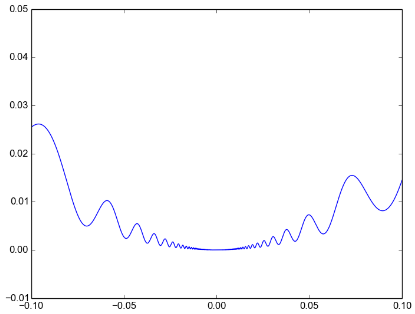

FIGURE 1. A close-up view of the function $f(x) = 2x^2 + x^2\sin(1/x)$ near 0. The point $x = 0$ is the unique global minimum, but is also an accumulation point of isolated local minima.

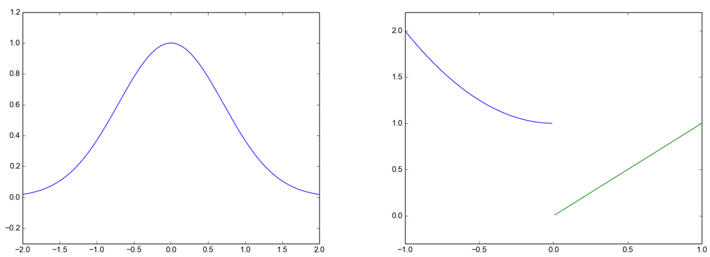

FIGURE 2. *Left:* The function $f(x) = e^{-x^2}$ obviously does not attain its minimum, because of the drop-off of the function values near infinity. *Right:* The existence of a minimiser of the function f defined by $f(x) = x^2 + 1$ for $x < 0$ and $f(x) = x$ for $x > 0$ depends on its value at 0. If $f(0) \le 0$, the point $x = 0$ is the unique global and local minimum. If, however, $f(0) > 0$, the function does not attain its minimum.

$x \le 0$. It turns out, however, that these two examples are in a sense the typical counter-examples to the existence of minimisers: In the case of the function e^{-x^2}, the problem is that the function to be minimised becomes smaller as the argument increases. In the case of the other counter-example, the problem is a discontinuity at the point where we would "naturally" expect the minimum. By excluding these two possibilities, that is, by requiring the function f to be continuous and to grow as its argument tends to infinity, we can indeed guarantee the existence of a minimiser. Because discontinuous functions can be important in some applications, it makes sense to try to obtain results for this type of functions as well, though.

For the following definition, recall that the lower limit of a sequence of real numbers $(z_k)_{k\in\mathbb{N}}$ is defined as

$$\liminf_{k\to\infty} z_k := \lim_{k\to\infty} \inf_{\ell\ge k} z_\ell.$$

This is equivalent to defining $\liminf_{k\to\infty} z_k$ as the smallest possible limit of convergent subsequences of z_k. Equivalently, $\liminf_{k\to\infty}$ is the infimum of all accumulation points of the sequence $(z_k)_{k\in\mathbb{N}}$ in the extended real line $\mathbb{R}\cup\{\pm\infty\}$.

Example 1.9. Consider the sequence $(z_k)_{k\in\mathbb{N}}$,

$$z_k = (-1)^k - \frac{1}{k}.$$

This sequence has the accumulation points -1 and $+1$. The lower limit of the sequence $(z_k)_{k\in\mathbb{N}}$ is the smaller value of these, that is, $\liminf_{k\to\infty} z_k = -1$.

Example 1.10. Consider the sequence $(z_k)_{k\in\mathbb{N}}$, $z_k = \sin(k)$. Every value in the interval $[-1,1]$ is an accumulation point of this sequence. Thus $\liminf_{k\to\infty} z_k = -1$.

Lemma 1.11. *The lower limit has the following properties:*

(1) The lower limit of a sequence $(z_k)_{k\in\mathbb{N}}$ always exists (in $\mathbb{R}\cup\{\pm\infty\}$).
(2) If the sequence $(z_k)_{k\in\mathbb{N}}$ converges, then $\liminf_{k\to\infty} z_k = \lim_{k\to\infty} z_k$.
(3) If $(y_k)_{k\in\mathbb{N}}$, $(z_k)_{k\in\mathbb{N}}$ are two sequences, then

$$\liminf_{k\to\infty}(y_k+z_k) \ge \liminf_{k\to\infty} y_k + \liminf_{k\to\infty} z_k.$$

Here we define $+\infty + (-\infty) := -\infty$.
(4) If $(y_k)_{k\in\mathbb{N}}$ is a sequence and $\lambda \ge 0$, then

$$\liminf_{k\to\infty} \lambda y_k = \lambda \liminf_{k\to\infty} y_k.$$

Here $\lambda(\pm\infty) = \pm\infty$ for $\lambda > 0$, and $0\cdot(\pm\infty) := 0$.
(5) If $(y_k^{(i)})_{k\in\mathbb{N}}$, $i \in I$, is a family of sequences (with an arbitrary index set I), then

$$\liminf_{k\to\infty} \sup_{i\in I} y_k^{(i)} \ge \sup_{i\in I} \liminf_{k\to\infty} y_k^{(i)}.$$

Proof. Exercise! □

Definition 1.12 (Lower semi-continuity)**.** A function $f\colon \mathbb{R}^d \to \mathbb{R}$ is called *lower semi-continuous*, if for every $x \in \mathbb{R}^d$ and every sequence $(x_k)_{k\in\mathbb{N}}$ converging to x we have

$$f(x) \le \liminf_{k\to\infty} f(x_k).$$

This means that, whenever we have a sequence x_k converging to x, the sequence of values $f(x_k)$ cannot have a limit that is smaller than $f(x)$. For instance:

- Every continuous function is lower semi-continuous.
- The function $f\colon \mathbb{R}\to\mathbb{R}$ defined by
$$f(x) = \begin{cases} x & \text{if } x > 0,\\ x^2+1 & \text{if } x \le 0,\end{cases}$$
is *not* lower semi-continuous.
- The function $f\colon \mathbb{R}\to\mathbb{R}$ defined by
$$f(x) = \begin{cases} x & \text{if } x \ge 0,\\ x^2+1 & \text{if } x < 0,\end{cases}$$
is lower semi-continuous.
- The function $f\colon \mathbb{R}\to\mathbb{R}$ defined by
$$f(x) = \begin{cases} \sin(1/x) & \text{if } x \ne 0,\\ -1 & \text{if } x = 0,\end{cases}$$
is lower semi-continuous.

- The function $f\colon \mathbb{R} \to \mathbb{R}$ defined by
$$f(x) = \begin{cases} +1 & \text{if } x > 0, \\ 0 & \text{if } x = 0, \\ -1 & \text{if } x < 0, \end{cases}$$
is *not* lower semi-continuous.

Lower semi-continuous functions play a prominent role in optimisation: On the one hand, they are amenable to minimisation, as the existence results in the following section show. On the other hand, they appear naturally in certain branches of optimisation in the form of min-max (or inf-sup) problems, that is, problems of the form
$$\inf_{x\in\Omega} \sup_{y\in W} g(x,y).$$
The following result shows that the function $h(x) := \sup_{y\in W} g(x,y)$ is lower semi-continuous (but not necessarily continuous), provided the function g is continuous in x for each y.

Lemma 1.13. *If $f_i\colon \mathbb{R}^d \to \mathbb{R}$, $i \in I$, is* any *family of continuous functions, then the function*
$$f(x) := \sup_{i\in I} f_i(x)$$
is lower semi-continuous.

Note that we do not require that the family in the previous result is finite or countable!

Proof. Let $x \in \mathbb{R}^d$ and assume that the sequence $(x_k)_{k\in\mathbb{N}}$ converges to x. Then
$$\liminf_{k\to\infty} f(x_k) = \liminf_{k\to\infty} \sup_{i\in I} f_i(x_k) \ge \sup_{i\in I} \liminf_{k\to\infty} f_i(x_k) = \sup_{i\in I} f_i(x) = f(x).$$
Thus f is lower semi-continuous at x. □

Remark 1.14. An alternative (equivalent) definition of lower semi-continuity is the following:

A function $f\colon \mathbb{R}^d \to \mathbb{R}$ is lower semi-continuous, if the *lower level set*
$$\Omega_\alpha := \bigl\{x \in \mathbb{R}^d : f(x) \le \alpha\bigr\}$$
is closed for every $\alpha \in \mathbb{R}$.

In other words: Whenever $\alpha \in \mathbb{R}$ and $(x_k)_{k\in\mathbb{N}}$ is a sequence that converges to some $x \in \mathbb{R}^d$ and $x_k \in \Omega_\alpha$ for all k (that is, $f(x_k) \le \alpha$), we have that $x \in \Omega_\alpha$ (that is, $f(x) \le \alpha$). Because this definition does not rely directly on sequences but rather on the notion of closedness, it can, in some situations, be less cumbersome to handle.

A further characterisation of lower semi-continuity can be obtained by considering convergence of points on the graph of f. One can show that a function $f\colon \mathbb{R}^d \to \mathbb{R}$ is lower semi-continuous, if and only if for every $x \in \mathbb{R}^d$ and every sequence $(x_k)_{k\in\mathbb{N}}$ such that the values $f(x_k)$ converge in the extended real line $\mathbb{R} \cup \{\pm\infty\}$ we have that
$$f(x) \le \lim_{k\to\infty} f(x_k). \tag{1}$$
In that sense, the difference between continuity and lower semi-continuity is that the equal-sign in the definition of continuity is replaced by a lower-than-or-equal sign. Be careful with this characterisation, though, as it is not sufficient that (1) holds for sequences $(x_k)_{k\in\mathbb{N}}$ where $f(x_k)$ converge to a *finite* limit, but also for those where $f(x_k)$ converges to an *infinite* limit value (specifically to $-\infty$). As an example, the function $f(x) = 1/x$ for $x \ne 0$ with $f(0) = 0$ is *not* lower semi-continuous.

We now discuss the existence of solutions of optimisation problems. The main result here is the following Lemma, which is a modification of the Weierstraß theorem from basic calculus.

Proposition 1.15. *Assume that* $f\colon \mathbb{R}^d \to \mathbb{R}$ *is lower semi-continuous and that* $\Omega \subset \mathbb{R}^d$ *is compact. Then the optimisation problem*

$$\min_{x\in\Omega} f(x)$$

admits a solution.

Proof. Denote

$$f^* := \inf_{x\in\Omega} f(x).$$

Then there exists a sequence $(x_k)_{k\in\mathbb{N}} \subset \Omega$ such that

$$\lim_{k\to\infty} f(x_k) = f^*.$$

Because of the compactness of Ω, the sequence $(x_k)_{k\in\mathbb{N}}$ has a subsequence $(x_{k'})_{k'}$ converging to some $x^* \in \Omega$. From the lower semi-continuity of f we now obtain that

$$f(x^*) \le \liminf_{k'\to\infty} f(x_{k'}) = f^* = \inf_{x\in\Omega} f(x).$$

This shows that x^* actually is a global minimum of f in Ω. ☐

This result covers the case of constrained optimisation problems where the feasible set Ω is compact. However, it also can be used in order to prove existence of solutions in more general settings as long as f is lower semi-continuous. The basic idea there is to show that all possible minimisers can only ever be found in some compact set, which makes it possible to reduce the problem to one with compact constraints. One possibility for this approach is shown in the following.

Definition 1.16. A function $f\colon \mathbb{R}^d \to \mathbb{R}$ is called coercive, if we have for every sequence $(x_k)_{k\in\mathbb{N}}$ with $\|x_k\| \to \infty$ that $f(x_k) \to \infty$.

That is, a function is coercive, if its value approaches infinity provided its arguments do (in the norm).

Example 1.17. Consider the following functions:

- The function $f\colon \mathbb{R} \to \mathbb{R}$, $f(x) = 1 - e^{-x^2}$, is not coercive, as $\lim_{x\to\pm\infty} f(x) = 1 \neq \infty$.
- The function $f\colon \mathbb{R} \to \mathbb{R}$, $f(x) = e^x$, is not coercive: Although we have $\lim_{x\to+\infty} f(x) = +\infty$, we also have that $\lim_{x\to-\infty} f(x) = 0$.
- Assume that $f\colon \mathbb{R} \to \mathbb{R}$, $f(x) = \sum_{k=0}^{n} c_k x^k$ with $c_n \neq 0$ is a polynomial of degree (precisely!) n. Then f is coercive, if and only if n is even and $c_n > 0$.
- Let $A \in \mathbb{R}^{d\times d}$ be a symmetric matrix and define $f\colon \mathbb{R}^d \to \mathbb{R}$, $f(x) = \langle x, Ax\rangle = x^T Ax$. Then f is coercive, if and only if A is positive definite.

Coercivity is important when we want to minimise a function, as this property allows us to exclude all points x with sufficiently large norm *a-priori*. This makes it possible to formulate existence theorems also for the case of unconstrained setting.

Theorem 1.18. *Assume that the function* $f\colon \mathbb{R}^d \to \mathbb{R}$ *is lower semi-continuous and coercive. Then the optimisation problem* $\min_{x\in\mathbb{R}^d} f(x)$ *admits at least one global minimiser* x^**.*

Proof. By assumption, the function f is coercive, which implies that there exists some $R > 0$ such that

$$f(x) > f(0) \qquad \text{whenever } \|x\| > R.$$

As a consequence, if the problem $\min_{x \in \mathbb{R}^d} f(x)$ admits a solution, this solution necessarily has to be contained in the ball $B_R := \{x \in \mathbb{R}^d : \|x\| \le R\}$, as every point $x \notin B_R$ is minorised by 0. This, however, implies that the problem $\min_{x \in \mathbb{R}^d} f(x)$ is equivalent to the constrained optimisation problem $\min_{x \in B_R} f(x)$. From Proposition 1.15 we now obtain the existence of a solution of the latter problem, which, because of the equivalence of the two problems, immediately implies the existence of the original, unconstrained one. □

While these existence results cover a large family of optimisation problems, they still might require some adaptation to concrete cases. Thus it is more important to understand the idea of the proof than the results themselves.

A general approach to proving existence of solutions of optimisation problems $\min_{x \in \Omega} f(x)$ with lower semi-continuous f is the following (also known as *the direct method in the calculus of variation*):

- Choose a minimising sequence, that is, a sequence $(x_k)_{k \in \mathbb{N}}$ such that $f(x_k)$ converges to $f^* := \inf_{x \in \Omega} f(x)$. Such a sequence always exists (provided that Ω is non-empty).
- Show by whatever means that you can choose the minimising sequence in such a way that it is bounded.
- Use the Weierstraß Theorem (which states that every bounded sequence in $\mathbb{R}^d$ has a convergent subsequence) to extract a convergent subsequence $(x_{k'})$.
- Show that the limit $x^* := \lim_{k'} x_{k'}$ satisfies $x^* \in \Omega$.
- Conclude from the lower semi-continuity of f that x^* solves the optimisation problem at hand.

A slightly less general, but simpler, approach is the following:

- Show by whatever means that all the possible candidates for a solution of the problem $\min_{x \in \Omega} f(x)$ lie in a closed and bounded subset X of Ω.
- Conclude from this that the problem $\min_{x \in \Omega} f(x)$ is equivalent to the problem $\min_{x \in X} f(x)$ in the sense that these problems have precisely the same global solutions.
- Use Proposition 1.15 to show that the problem $\min_{x \in X} f(x)$ admits a solution and conclude that the original problem does so as well.

2. Characterisation of Solutions

We now discuss how we can characterise local solutions of unconstrained optimisation problems by the properties of their derivatives. Later in the course, we will look at the much harder problem of characterising solutions of *constrained* problems as well. The basis of everything that follows are first and second order Taylor expansions of sufficiently regular functions.

To that end, recall that $\nabla f(x) \in \mathbb{R}^d$ denotes the gradient of the differentiable function $f\colon \mathbb{R}^d \to \mathbb{R}$ (that is, the vector of its partial derivatives), and $H_f(x) \in \mathbb{R}^{d \times d}$ denotes the Hessian of the twice differentiable function $f\colon \mathbb{R}^d \to \mathbb{R}$, that is, the matrix of its second order partial derivatives.

Theorem 2.1. *Assume that $f \in C^1(\mathbb{R}^d)$ and that $x^* \in \mathbb{R}^d$. Then*

$$f(x) = f(x^*) + \langle \nabla f(x^*), x - x^* \rangle + o(\|x - x^*\|) \qquad \text{as } x \to x^*.$$

If $f \in C^2(\mathbb{R}^d)$, then

$$f(x) = f(x^*)+\langle \nabla f(x^*), x-x^*\rangle+\frac{1}{2}\langle x-x^*, H_f(x^*)(x-x^*)\rangle+o(\|x-x^*\|^2) \text{ as } x \to x^*.$$

Moreover there exists for every $x \in \mathbb{R}^d$ some $0 < t_x < 1$ such that

$$f(x) = f(x^*) + \langle \nabla f(x^*), x - x^*\rangle + \frac{1}{2}\langle x - x^*, H_f(tx + (1 - t)x^*)(x - x^*)\rangle.$$

Proof. See Calculus 2. □

2.1. **Necessary conditions.**

Theorem 2.2 (First order necessary condition)**.** *Assume that $f \in C^1(\mathbb{R}^d)$ and that x^* is a local solution of the optimisation problem*

$$\min_{x\in\mathbb{R}^d} f(x).$$

Then

$$\nabla f(x^*) = 0.$$

Proof. Assume to the contrary that $\nabla f(x^*) =: p \neq 0$. Then a first order Taylor series expansion of $f(x^* - tp)$ with $t > 0$ implies that

$$\lim_{t\to 0^+} \frac{f(x^* - tp) - f(x^*)}{t} = \lim_{t\to 0^+} \frac{\langle \nabla f(x^*), -tp\rangle + o(\|tp\|)}{t} = -\|p\|^2 + \lim_{t\to 0^+} \frac{o(\|tp\|)}{t} = -\|p\|^2 < 0.$$

This, however, means that $f(x^* - tp) < f(x^*)$ for all sufficiently small $t > 0$, which in turn implies that x^* cannot be a local minimiser of f. □

Theorem 2.3 (Second order necessary conditions)**.** *Assume that $f \in C^2(\mathbb{R}^d)$ and that x^* is a local solution of the optimisation problem*

$$\min_{x\in\mathbb{R}^d} f(x).$$

Then

$$\nabla f(x^*) = 0$$

and $H_f(x^*) \in \mathbb{R}^{d\times d}$ is positive semi-definite.

Proof. From Theorem 2.2 we already know that $\nabla f(x^*) = 0$. Assume therefore that $\nabla f(x^*) = 0$ but $H_f(x^*)$ is not positive semi-definite. Then $H_f(x^*)$ has a negative eigenvalue $\lambda < 0$ with associated eigenvector $p \in \mathbb{R}^d$. Now a second order Taylor series expansion implies that

$$\lim_{\substack{t\to 0\\ t\neq 0}} \frac{f(x^* + tp) - f(x^*)}{t^2} = \lim_{\substack{t\to 0\\ t\neq 0}} \frac{\langle \nabla f(x^*), tp\rangle + \langle tp, H_f(x^*)(tp)\rangle/2 + o(\|tp\|^2)}{t^2}$$

$$= \lim_{\substack{t\to 0\\ t\neq 0}} \frac{t^2\langle p, H_f(x^*)p\rangle + o(\|tp\|^2)}{t^2} = \lim_{\substack{t\to 0\\ t\neq 0}} \frac{t^2\lambda\|p\|^2 + o(\|tp\|^2)}{t^2} = \lambda\|p\|^2 < 0.$$

Thus $f(x^* + tp) < f(x^*)$ for all $t \neq 0$ sufficiently close to 0, which again implies that x^* cannot be a local minimiser of f. □

Note that the previous conditions are necessary, but not sufficient, for x^* to be a local minimiser of f:

Example 2.4. Consider the function $f : \mathbb{R} \to \mathbb{R}$, $f(x) = -x^4$. Then $f'(0) = 0$ and $f''(0) = 0$. Thus the second order necessary optimality conditions are satisfied for $x^* = 0$, but 0 is no local minimser of f (but rather a global maximiser).

2.2. **Sufficient conditions.** In order to obtain a sufficient optimality condition, one has to strengthen the positive semi-definiteness of the Hessian to positive definiteness:

Theorem 2.5. *Assume that $f \in C^2(\mathbb{R}^d)$ and that $x^* \in \mathbb{R}^d$ is such that $\nabla f(x^*) = 0$ and $H_f(x^*)$ is positive definite. Then x^* is an isolated and strict local minimiser of f.*

Proof. Since $H_f(x^*)$ is positive definite, there exists $\sigma > 0$ (the smallest singular value of $H_f(x^*)$) such that

$$\langle p, H_f(x^*)\rangle \geq \sigma \|p\|^2$$

for all $p \in \mathbb{R}^d$. Thus we obtain from a second order Taylor series expansion of f around x^* that

$$\begin{aligned} f(x^* + p) &= f(x^*) + \langle \nabla f(x^*), p\rangle + \frac{1}{2}\langle p, H_f(x^*)p\rangle + o(\|p\|^2) \\ &\geq f(x^*) + \frac{\sigma}{2}\|p\|^2 + o(\|p\|^2) \text{ as } p \to 0. \end{aligned}$$

For sufficiently small $\|p\|$, the term $o(\|p\|^2)$ is dominated by $\frac{\sigma}{4}\|p\|^2$. Thus we have that

$$f(x^* + p) \geq f(x^*) + \frac{\sigma}{4}\|p\|^2$$

for all p with $\|p\|$ sufficiently small. This shows that x^* is a strict local minimiser of f.

We still have to show that x^* is an isolated local minimiser of f. To that end, we perform a first order Taylor series expansion of ∇f around x^* and obtain that

$$\nabla f(x^* + p) = \nabla f(x^*) + H_f(x^*)p + o(\|p\|) = H_f(x^*)p + o(\|p\|).$$

Taking the inner product with p, this implies that

$$\langle p, \nabla f(x^* + p)\rangle = \langle p, H_f(x^*)\rangle + \langle p, o(\|p\|)\rangle \geq \sigma\|p\|^2 + o(\|p\|^2).$$

Again, the term $\sigma\|p\|^2/2$ dominates the term $o(\|p\|^2)$ for sufficiently small $\|p\|$, which implies that

$$\langle p, \nabla f(x^* + p)\rangle \geq \frac{\sigma}{2}\|p\|^2$$

for all p with $\|p\|$ sufficiently small. In particular, this shows that $\nabla f(x^* + p) \neq 0$ if $p \neq 0$ is sufficiently close to 0. In view of the first order necessary optimality condition, this further implies that no point $x^* + p$ with $p \neq 0$ sufficiently close to 0 can be a local minimiser of f. In other words, x^* is an isolated local minimiser. □

Definition 2.6. Let $f \in C^1(\mathbb{R}^d)$. A point $x^* \in \mathbb{R}^d$ with $\nabla f(x^*) = 0$ is called a *critical point* of f.

Remark 2.7. From the first order optimality condition, it follows that every local minimiser (or maximiser) of a C^1 function f is also a critical point. Conversely, every critical point may be a local minimiser, a local maximiser, or neither of these. Sometimes, the points in the latter category, that is, critical points that are neither local minimisers nor local maximisers, are called saddle points of f. We will, however, not adopt this notation and rather use the term "saddle point" in a more restrictive setting later in the course.

Example 2.8. Define $f\colon \mathbb{R}^2 \to \mathbb{R}$,

$$f(x, y) = 3x^4 + 4x^3 + 12y^2 - 24xy.$$

We want to compute the local and global minimisers of f on $\mathbb{R}^2$. To that end, we first compute the gradient and Hessian of f, which are

$$\nabla f(x,y) = \begin{pmatrix} 12x^3 + 12x^2 - 24y \\ 24y - 24x \end{pmatrix}$$

and

$$H_f(x,y) = \begin{pmatrix} 36x^2 + 24x & -24 \\ -24 & 24 \end{pmatrix}.$$

For the computation of the critical points of f, we have to solve the system of equations

$$\nabla f(x,y) = \begin{pmatrix} 12x^3 + 12x^2 - 24y \\ 24y - 24x \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

From the second equation, we immediately get that $x = y$. Inserting this result in the first equation, we obtain the condition

$$12x^3 + 12x^2 - 24x = 0.$$

This equation has three solutions $x = -2$, $x = 0$, and $x = 1$. We thus have the three critical points $(-2,-2)$, $(0,0)$, and $(1,1)$.

- The point $(-2,-2)$. Here the Hessian of f is

$$H_f(-2,-2) = \begin{pmatrix} 96 & -24 \\ -24 & 24 \end{pmatrix},$$

  which is positive definite. Thus this point is a strict local minimiser.
- The point $(0,0)$. Here the Hessian of f is

$$H_f(0,0) = \begin{pmatrix} 0 & -24 \\ -24 & 24 \end{pmatrix},$$

  which is indefinite. Thus this point is no local minimiser (and neither a maximiser).
- The point $(1,1)$. Here the Hessian of f is

$$H_f(1,1) = \begin{pmatrix} 60 & -24 \\ -24 & 24 \end{pmatrix},$$

  which is positive definite. Thus this point is a strict local minimiser.

Finally, we want to determine whether any of the points $(-2,-2)$ and $(1,1)$ is a global minimiser. Computing the function values at these points, we obtain that $f(-2,-2) = -8$ and $f(1,1) = -5$. Thus $(-2,-2)$ is the only possible candidate for a global minimiser. Note, however, that we cannot yet conclude that $(-2,-2)$ actually *is* a global minimiser: From all what we know about the function, it could still be possible that f has no global minimisers at all and, for instance, be unbounded below. We thus try to investigate whether the function f is coercive. To that end, we note that the inequality

$$4xy \le 4x^2 + y^2$$

holds for all x, $y \in \mathbb{R}$: This follows from the fact that

$$4x^2 - 4xy + y^2 = (2x - y)^2 \ge 0.$$

Thus we can estimate

$$f(x,y) = 3x^4 + 4x^3 + 12y^2 - 24xy \ge 3x^4 + 4x^3 - 24x^2 + 6y^2 \ge x^4 + 6y^2 - C$$

for some $C > 0$ and all $x, y \in \mathbb{R}$.[1] This shows that f is coercive. Since it is obviously continuous, we can conclude that f admits a global minimum, and, as seen above, the only candidate for this global minimum is the point $(-2, -2)$. This shows that $(-2, -2)$ is actually the (unique) strict global minimum of f.

3. Optimisation and Convex Functions

3.1. **Convex Functions and Sets.**

Definition 3.1 (Convex Function)**.** A function $f : \mathbb{R}^d \to \mathbb{R}$ is *convex*, if for every $x, y \in \mathbb{R}^d$ and $0 \le \lambda \le 1$ the inequality

$$f\big(\lambda x + (1 - \lambda)y\big) \le \lambda f(x) + (1 - \lambda)f(y)$$

holds.

The function f is *strictly convex*, if for every $x \neq y \in \mathbb{R}^d$ and $0 < \lambda < 1$ the inequality

$$f\big(\lambda x + (1 - \lambda)y\big) < \lambda f(x) + (1 - \lambda)f(y)$$

holds. That is, the inequality defining the convexity of a function is strict whenever possible.

More graphically, this means that for each pair of points $(x, f(x))$ and $(y, f(y))$ lying on the graph of f, the connecting line segment remains above (or rather: not below) the graph. It is strictly convex if the connecting line segment stays strictly above the graph. See Figure 3.

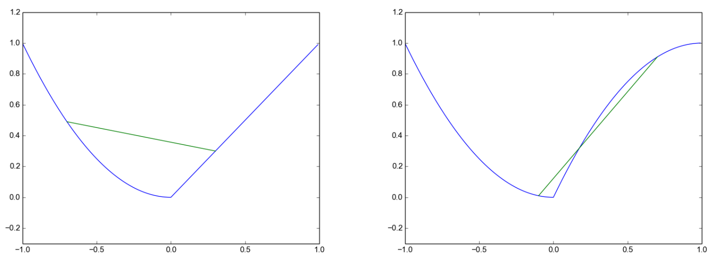

Figure 3. *Left:* Typical example of a convex (but not strictly convex) function. Note that no differentiability is assumed. *Right:* Typical example of a non-convex function. There exist points on the graph such that the connecting line segment does not lie completely above the graph.

Similarly, we can also define convex sets:

Definition 3.2. A set $C \subset \mathbb{R}^d$ is convex, if for all points $x, y \in C$ and $0 \le \lambda \le 1$ we have

$$\lambda x + (1 - \lambda)y \in C.$$

That is, a set is convex, if whenever we are given two points x and y in C the whole line segment connecting these two points is also contained in C.

An important property of convex functions, which we will need later in the course when dealing with constrained optimisation problems, is the fact that lower level sets of convex functions are convex sets:

[1]If we really want, we can, for instance, argue here that $2x^4 + 4x^3 - 24x^2 \ge 0$ whenever $|x| \ge 5$ and that $2x^4 + 4x^3 - 24x^2 \ge -4 \cdot 5^3 - 24 \cdot 5^2 = -1100$ whenever $|x| \le 5$. Thus $C = -1100$ would work.

Lemma 3.3. *Given a function $f\colon \mathbb{R}^d \to \mathbb{R}$, we define for $C \in \mathbb{R}$ the C-lower level set $S_C(f)$ of f as*

$$S_C(f) := \{x \in \mathbb{R}^d : f(x) \le C\}.$$

Assume now that $f\colon \mathbb{R}^d \to \mathbb{R}$ is a convex function. Then $S_C(f) \subset \mathbb{R}^d$ is a convex set for every $C \in \mathbb{R}$.

Proof. Let $C \in \mathbb{R}$ be fixed and assume that $x, y \in S_C(f)$. That is, $f(x) \le C$ and $f(y) \le C$. Let moreover $0 < \lambda < 1$. Then the convexity of f implies that

$$f\big(\lambda x + (1-\lambda)y\big) \le \lambda f(x) + (1-\lambda) f(y) \le \lambda C + (1-\lambda)C = C,$$

showing that $\lambda x + (1-\lambda)y \in S_C(f)$ which in turn shows the convexity of $S_C(f)$. □

Corollary 3.4. *Assume that $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex. Then the set of global minimisers of f is convex.*

Proof. This is a direct consequence of Lemma 3.3 with $C := \inf_{x\in\mathbb{R}^d} f(x)$. □

Remark 3.5. There is a very close connection between convex sets and convex functions: One can show that a function $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex, if and only if the so-called *epigraph* of f, which is the subset of $\mathbb{R}^d \times \mathbb{R}$ consisting of all points (x,t) with $t \ge f(x)$, is a convex set.

It is easy to show the following properties of convex functions:

- If the functions f, $g\colon \mathbb{R}^d \to \mathbb{R}$ are convex, then so is the function $f + g$.
- If $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex and $\lambda \ge 0$, then also the function λf is convex.
- Every linear (or affine) function is convex.
- If both f and $-f$ are convex, then the function f is affine (that is, $f(x) = \langle a, x\rangle + b$ for some $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$).
- If f and g are convex functions, then the function h defined by $h(x) := \max\{f(x), g(x)\}$ is also convex.

From the point of view of optimisation, one of the many good properties of convex functions is the fact that there is no difference between global and local minima; instead, every local minimum is automatically global, as the following result shows.

Lemma 3.6. *Assume that $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex. Then every local minimiser of f is already a global minimiser.*

Proof. Assume to the contrary that $x \in \mathbb{R}^d$ is no global minimiser of f. Then there exists $y \ne x$ with $f(y) < f(x)$. However, because of the convexity of f we have for every $0 < \lambda < 1$ that

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

That is, setting $x_\lambda = \lambda x + (1-\lambda)y$ we have $f(x_\lambda) < f(x)$ for all $0 < \lambda < 1$. Since $x_\lambda \to x$ as $\lambda \to 1$, this shows that x cannot be a local minimiser of f. □

3.2. **Differentiable Convex Functions.** In the definition of convex functions above, we have not assumed any regularity of f (apart from f only taking finite values). Indeed, one of the main advantages of the (rather extensive) theory of convex functions is that it allows to deal with non-differentiable functions using almost the same methods as we would use for differentiable functions. In particular, it is possible to introduce generalised notions of derivatives that in turn can be used for the characterisation and computation of solutions of optimisation problems. However, we will consider in the following *differentiable* convex functions, and we will study what the convexity of a function implies for its derivative.[2]

[2] We will not discuss non-continuous convex functions—the main reason being that they do not exist in the setting used here: It can be shown that every function $f\colon \mathbb{R}^d \to \mathbb{R}$ is not only continuous, but actually locally Lipschitz continuous. Moreover, this implies that convex functions are

Proposition 3.7. *Assume that the function $f\colon \mathbb{R}^d \to \mathbb{R}$ is differentiable. Then f is convex, if and only if for every x, $y \in \mathbb{R}^d$ the inequality*

$$f(y) \geq f(x) + \langle \nabla f(x), y-x\rangle \tag{2}$$

is satisfied.

Proof. Assume first that f is convex and let $x \neq y \in \mathbb{R}^d$. We note that $\langle \nabla f(x), y-x\rangle$ is precisely the directional derivative of f at the point x in direction $(y-x)$, that is,

$$\langle \nabla f(x), y-x\rangle = \lim_{t\to 0} \frac{1}{t}\big(f(x+t(y-x)) - f(x)\big).$$

Using the convexity of f and that

$$x + t(y-x) = ty + (1-t)x,$$

we can estimate the right hand side by

$$\frac{1}{t}\big(f(x+t(y-x)) - f(x)\big) \leq \frac{1}{t}\big(tf(y) + (1-t)f(x) - f(x)\big) = f(y) - f(x)$$

for $0 < t < 1$, implying that

$$\langle \nabla f(x), y-x\rangle \leq f(y) - f(x),$$

which is precisely (2).

Assume now that the inequality (2) holds for all x, $y \in \mathbb{R}^d$. Let w, $z \in \mathbb{R}^d$ and $0 \leq \lambda \leq 1$. Denote moreover

$$x := \lambda w + (1-\lambda)z.$$

Then the inequality (2) implies that

$$\begin{aligned} f(w) &\geq f(x) + \langle \nabla f(x), w-x\rangle,\\ f(z) &\geq f(x) + \langle \nabla f(x), z-x\rangle. \end{aligned} \tag{3}$$

Note moreover that

$$w - x = (1-\lambda)(w-z) \qquad \text{and} \qquad z - x = \lambda(z-w).$$

Thus, if we multiply the first line in (3) with λ, the second line with $1-\lambda$, and then add the two inequalities, we obtain

$$\begin{aligned} \lambda f(w) + (1-\lambda)f(z) \geq f(x) + \lambda\langle \nabla f(x), (1-\lambda)(w-z)\rangle + (1-\lambda)\langle \nabla f(x), \lambda(z-w)\rangle \\ = f\big(\lambda w + (1-\lambda)z\big). \end{aligned}$$

Since w and z were arbitrary, this proves the convexity of f. □

Remark 3.8. Following basically the same proof as above and strategically replacing inequalities by strict inequalities, one can show that a differentiable function f is strictly convex, if and only if

$$f(y) > f(x) + \langle \nabla f(x), y-x\rangle$$

whenever $x \neq y \in \mathbb{R}^d$.

---

actually almost everywhere differentiable (in the sense that the set of points where the derivative does not exists has Lebesgue measure zero).

It *is* possible to construct non-continuous convex functions, but only if one either restricts the domain of the function to some (non-open) convex subset C of $\mathbb{R}^d$ (that is, we have a function $f\colon C \to \mathbb{R}$), or one allows the function to take the value $+\infty$. For instance, the function $f(x) = x$ for $x > 0$, and $f(0) = 1$ is convex on the half line $\mathbb{R}_{\geq 0}$, but discontinuous at 0.

Remark 3.9. Originally, we have introduced the convexity of functions by requiring that the graph of the function lies below all the line segments connecting two points on the graph. From the previous result, however, we obtain an alternative characterisation of convexity that requires that the graph of the function lies *above* all the tangents to the function. While the first definition of convexity may seem to be more natural and more general, a definition of convexity of functions by means of the properties of the "tangents" to the function

As an immediate consequence of Proposition 3.7 one obtains the result that the first order necessary condition for a minimiser is, in the case of convex functions, also a sufficient condition. More precisely, the following holds:

Corollary 3.10. *Assume that $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex and differentiable. Then x^* is a global minimiser of f, if and only if $\nabla f(x^*) = 0$.*

Proof. First recall that the condition $\nabla f(x^*) = 0$ is, independent of the convexity of f, a necessary condition for x^* to be a global (and indeed already local) minimiser. Thus we only need to show that this condition actually implies that x^* *is* a global minimiser. Assume therefore that $\nabla f(x^*) = 0$ and let $y \in \mathbb{R}^d$. Then Proposition 3.7 implies that

$$f(y) \geq f(x^*) + \langle \nabla f(x^*), y - x^* \rangle = f(x^*).$$

Thus x^* is a global minimiser. □

Proposition 3.11. *Assume that the function $f\colon \mathbb{R}^d \to \mathbb{R}$ is differentiable. Then f is convex, if and only if for every x, $y \in \mathbb{R}^d$ the inequality*

$$\langle \nabla f(y) - \nabla f(x), y - x \rangle > 0 \tag{4}$$

is satisfied.

Proof. Assume first that f is convex and that x, $y \in \mathbb{R}^d$. Then Proposition 3.7 implies the two inequalities

$$\begin{aligned} f(y) &\geq f(x) + \langle \nabla f(x), y - x \rangle, \\ f(x) &\geq f(y) + \langle \nabla f(y), x - y \rangle. \end{aligned}$$

Adding these inequalities, we obtain that

$$f(y) + f(x) \geq f(x) + f(y) + \langle \nabla f(x) - \nabla f(y), y - x \rangle,$$

which simplifies to

$$\langle \nabla f(y) - \nabla f(x), y - x \rangle \geq 0.$$

Conversely, assume that (4) holds for all x, $y \in \mathbb{R}^d$. Let moreover z, $x \in \mathbb{R}^d$. Then the mean value theorem implies that there exists $0 < \lambda < 1$ such that

$$f(z) - f(x) = \langle \nabla f\big(x + \lambda(z - x)\big), z - x \rangle.$$

Now we can write

$$\begin{aligned} \langle \nabla f\big(x + \lambda(z - x)\big), z - x \rangle &= \langle \nabla f\big(x + \lambda(z - x)\big) - \nabla f(x), z - x \rangle + \langle \nabla f(x), z - x \rangle \\ &= \frac{1}{\lambda} \langle \nabla f\big(x + \lambda(z - x)\big) - \nabla f(x), x + \lambda(z - x) - x \rangle + \langle \nabla f(x), z - x \rangle. \end{aligned}$$

Applying (4) with $y = x + \lambda(z - x)$ (and recalling that $\lambda > 0$), it follows that

$$\langle \nabla f\big(x + \lambda(z - x)\big), z - x \rangle \geq \langle \nabla f(x), z - x \rangle.$$

Thus

$$f(z) - f(x) \geq \langle \nabla f(x), z - x \rangle,$$

which is just a reformulation of (2). Since this holds for every z, $x \in \mathbb{R}^d$, we obtain the convexity of f from Proposition (3.7). □

Remark 3.12. A function $G\colon \mathbb{R}^d \to \mathbb{R}^d$ is called *monotone*, if

$$\langle G(y) - G(x), y - x\rangle \geq 0$$

for every y, $x \in \mathbb{R}^d$. (In the particular case $d = 1$, this actually is the same as stating that G is monotonically non-decreasing; do check this equivalence!) With this notation, Proposition 3.11 can be reformulated as stating that f is convex if and only if ∇f is monotone.

3.3. **Hessians of Convex Functions.**

Proposition 3.13. *A twice differentiable function $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex, if and only if the Hessian $H_f(x)$ is positive semi-definite for all $x \in \mathbb{R}^d$.*

Proof. Assume first that f is convex and let $x \in \mathbb{R}^d$. Define moreover the function $g\colon \mathbb{R}^d \to \mathbb{R}$ setting

$$g(y) := f(y) - \langle \nabla f(x), y - x\rangle.$$

Since the mapping $y \mapsto -\langle \nabla f(x), y - x\rangle$ is affine, it follows that g is convex. Moreover

$$\nabla g(y) = \nabla f(y) - \nabla f(x)$$

and

$$H_g(y) = H_f(y)$$

for all $y \in \mathbb{R}^d$. In particular, $\nabla g(x) = 0$. Thus Corollary 3.10 implies that x is a global minimiser of g. Now the second order necessary condition for a minimiser implies that $H_g(x)$ is positive semi-definite. Since $H_g(x) = H_f(x)$ and x was arbitrary, this proves that the Hessian of f is positive semi-definite for all $x \in \mathbb{R}^d$.

Now assume that the Hessian $H_f(x)$ of f is positive semi-definite for all $x \in \mathbb{R}^d$. Let moreover x, $y \in \mathbb{R}^d$. Then Taylor's theorem implies that

$$f(y) = f(x) + \langle \nabla f(x), y - x\rangle + \frac{1}{2}\langle y - x, H_f(x + t(y - x))(y - x)\rangle$$

for some $0 \leq t \leq 1$. Since H_f is everywhere positive semi-definite, the quadratic term in this equation is always non-negative. Thus we can estimate

$$f(y) \geq f(x) + \langle \nabla f(x), y - x\rangle.$$

Proposition 3.7 proves now the convexity of f. □

Remark 3.14. There is *some* relation between the strict convexity of a function f and the positive definiteness of its Hessian. However, this relation is not completely straight-forward. It is possible to show (and actually pretty simple to show) that a function $f\colon \mathbb{R}^d \to \mathbb{R}$ is strictly convex, if its Hessian $H_f(x)$ is positive definite for all x. However, the converse direction does not hold: The strict convexity of a function f does not imply that its Hessian is everywhere positive definite. As an example consider the function $f\colon \mathbb{R} \to \mathbb{R}$, $f(x) = x^4$. This function is strictly convex, but $f''(0) = 0$. Still, it is possible to characterise the strict convexity of a univariate function $f\colon \mathbb{R} \to \mathbb{R}$ by the condition that the set of points $x \in \mathbb{R}$ with $f''(x) > 0$ is dense. Thus a twice differentiable function $f\colon \mathbb{R} \to \mathbb{R}$ is strictly convex, if and only if the set $\{x \in \mathbb{R} : f''(x) > 0\}$ is dense in $\mathbb{R}$. To the best of my knowledge, there exists no (simple) generalisation of this characterisation to multivariate functions.

3.4. **Summary.** From the viewpoint of optimisation, the main results concerning convex functions (that we will need/refer to during this class) are:

- Convexity of a differentiable function can either characterised by the fact that all secants lie above the graph (Definition 3.1) or that all tangents lie below the graph (Proposition 3.7).
- If a function $f\colon \mathbb{R}^d \to \mathbb{R}$ is convex and differentiable, then the first order necessary condition for a minimum is actually sufficient. That is, the minimisation of f is equivalent to the solution of the equation
$$\nabla f(x) = 0.$$
- A function f is convex, if its Hessian is everywhere positive semi-definite. This allows us to test whether a given function is convex.
- If the Hessian of a function is everywhere positive definite, then the function is strictly convex. The converse does not hold.

Trondheim, Norway
Email address: `markus.grasmair@gmail.com`