ACM 41030 / ACM 40990

Chapter 1 - Intro to Optimization

Why? In many ML algorithms, it is required to solve an optimization problem. But these algorithms are now mature, and involve "only" running some code. But what if your code throws up an error — we should try to understand these error messages. Otherwise, we run the risk of GIGO.

The module(s) have two parts:

- Weeks 1–7: Continuous optimization
- Weeks 8–12:
  - ACM 40990: applications of continuous optimization to ML.
  - ACM 41030: We look at constrained optimization.

Wider context:

- Finance — portfolio management
- Manufacturing
- Operations Research

In nature, optimization is important:

- Systems tend to equilibrium — a state of minimum energy.
- Fermat's Principle of Least time — rays of light follow paths that minimize the time of travel.

Nature provides inspiration to create general optimization algorithms, applicable in many contexts — e.g. Simulated Annealing.

Terminology:

- $\underline{x} \in \mathbb{R}^n$ is the vector of variables, also called the unknowns or the parameters.
- f is the cost function (objective function):

$$f : \Omega \to \mathbb{R} \qquad (\Omega \subset \mathbb{R}^n)$$
$$\underline{x} \mapsto f(\underline{x})$$

The goal is to minimize f.

- c_i are the constraint functions. These are scalar-valued, and they define certain equalities or inequalities. The vector $\underline{x}$ has to satisfy these equalities / inequalities.

The fundamental Optimization Problem (OP):

Constrained :

$$\min_{\underline{x} \in \mathbb{R}^n} f(\underline{x}) \text{ subject to } \begin{cases} c_i(\underline{x}) = 0, & i \in \mathcal{E} \\ c_i(\underline{x}) \geq 0, & i \in \mathcal{I} \end{cases}$$

where $\mathcal{E}$ is a set of indices labelling the equality constraints and $\mathcal{I}$ is a set of indices labelling the inequality constraints.

Unconstrained :

$$\min_{\underline{x} \in \mathbb{R}^n} f(\underline{x})$$

Numerical algorithms for the unconstrained case:

- Line Search
- Newton
- Trust- Region

(+ Simulated Annealing)

Example : $\underbrace{\qquad\qquad\qquad\qquad}_{}$ $f(\underline{x})$, $\underline{x} = (x_1, x_2)^T$

$\overline{\underline{IV}}$

$$\min \quad \overbrace{(x_1-2)^2 + (x_2-1)^2}^{f(\underline{x})}$$

$$\text{subject to:} \quad \begin{cases} x_1^2 - x_2 \leq 0 \\ x_1 + x_2 \leq 2 \end{cases}$$

Constraint functions:

$$c_1(\underline{x}) = -x_1^2 + x_2, \qquad c_1(\underline{x}) \geq 0$$

$$c_2(\underline{x}) = -x_1 - x_2 + 2, \qquad c_2(\underline{x}) \geq 0.$$

$$\mathcal{E} = \emptyset, \quad \mathcal{I} = \{1, 2\}.$$

Solution — via a graphical argument.

Re-write the OP as:

$$f(x,y) = (x-2)^2 + (y-1)^2$$

c_1: $\quad x^2 - y \leq 0 \Rightarrow y \geq x^2$.

c_2: $\quad x + y \leq 2 \Rightarrow y \leq 2 - x$.

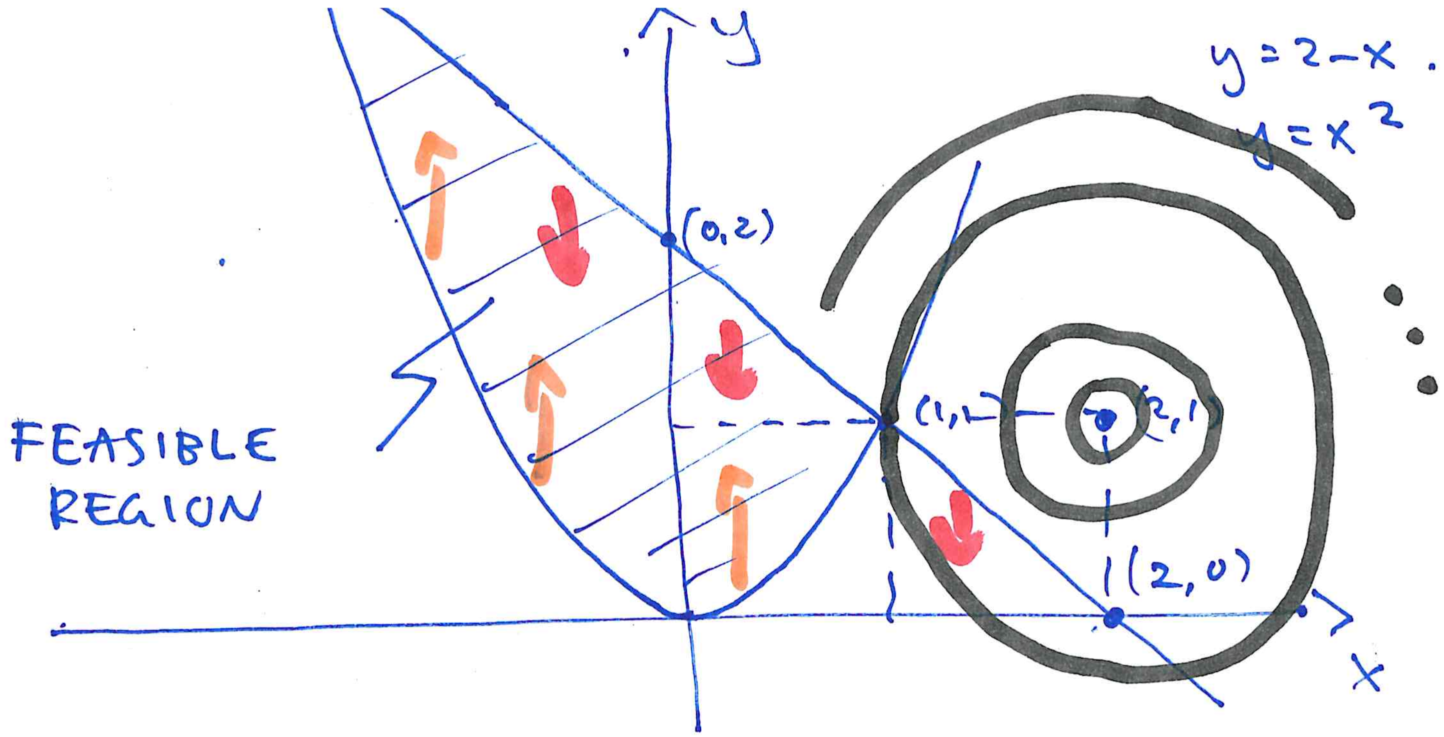

Cost function:

$$f(x,y) = (x-2)^2 + (y-1)^2 .$$

Contours of cost function:

$$f(x,y) = \text{Const.} = R^2 .$$

$$\Rightarrow (x-2)^2 + (y-1)^2 = R^2 .$$

– A circle of centre $(2,1)$, radius R.

To solve the OP, we take:

– R as small as possible

– while staying inside / on the boundary of the feasible region.

From the picture:

$$\underline{x}_* = (1,1)^T .$$

Definition: The feasible region is the set of all points satisfying the constraints.

For inequality constraints, the feasible region is the region enclosed by the constraint boundaries.

Notation: We use x_* to denote a solution of the OP.

Solution via pen-and-paper methods.
So why numerical optimization:

- 2D problem - can be done with pen and paper.
- For higher-dimensional problems, a computer is needed - hence NUMERICAL OPTIMIZATION.

"Housekeeping"

- Tuesday's lecture will be replaced with a pre-recorded lecture.
- But next Tuesday the Data and Computational Science class will use that slot for an info session with Dr. Sarp Akcay.
- Date and venue for midterm is now confirmed – email.

Now: we continue with Chapter 1, talking about the motivation and applications.

Example: A chemical company has the following setup:

- Two factories F_1 and F_2
- Twelve retail outlets $R_1, R_2, \dots, R_{12}$

Furthermore,

- Each factory F_i can produce a_i tons of product per week (capacity)
- Each retail outlet has a demand b_j tons of product per week.
- c_{ij} = cost of shipping 1 ton from factory i to retail outlet j
- x_{ij} = number of tons shipped from factory i to retail outlet j.

$$f = \sum_{i=1}^{2} \sum_{j=1}^{12} c_{ij} x_{ij} \qquad \text{COST FUNCTION}$$

Constraints:

- [illegible] : $\sum^{12} x_{ij} \leq a_i$

Constraints :

Capacity : $\sum_{j=1}^{12} x_{ij} \leq a_i$

Demand : $\sum_{i=1}^{2} x_{ij} \begin{cases} = b_j & \text{equality constraint} \\ & \text{OR} \\ \geq b_j & \text{inequality constraint} \end{cases}$

Positivity : $x_{ij} \geq 0$ for all $i \in \{1,2\}$, $j \in \{1,..,12\}$

The array x_{ij} can be reshaped to give a vector $\underline{x} \in \mathbb{R}^{24}$.

So the variable for the OP is $\underline{x} \in \mathbb{R}^{24}$.

The OP is linear in $\underline{x}$. Solution via linear programming.

Continuous versus discrete optimization:

- $\underline{x} \in \mathbb{R}^n \Rightarrow$ continuous optimization
- $\underline{x} \in \mathbb{N}^n$, $\underline{x} \in \{0,1\}^n$.

Example: Power grid, n different types of powerplant (pp) - e.g. wind, solar, gas, ...

$$\underline{x} = \left(\overset{pp1}{\cdots}, \overset{pp2}{\cdots}, \ldots, \overset{ppn}{\cdots} \right)$$

Thus, $\underline{x} \in \mathbb{N}_0^n$.

Example: Putting a factory in city i (y/n)

$$\underline{x} = \left(\overset{\text{city } 1}{}, \overset{\text{city } 2}{}, \ldots, \overset{\text{city } n}{} \right)$$

Thus, $\underline{x} \in \{0,1\}^n$

We focus in this module on continuous optimization.

Global versus local optimization:

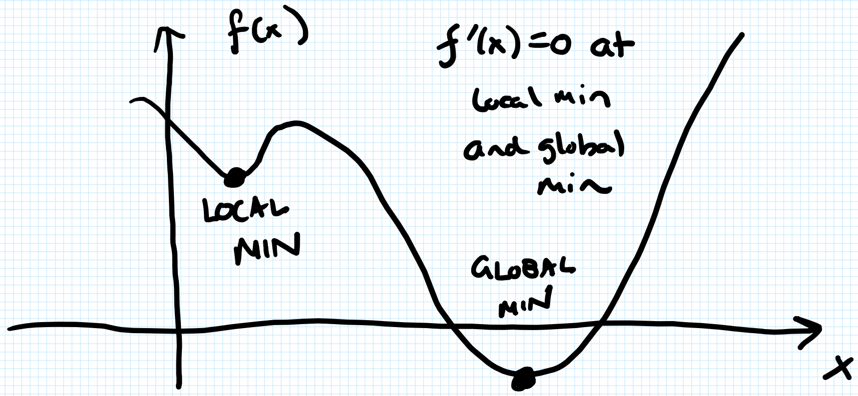

- Linesearch / Trust-Region / Newton will find a local min.
- Global optimization (Simulated Annealing) will find the global min.

For convex functions there is a unique minimum.

Convex Sets (§ 1.3)

Definition: Let $S \subset \mathbb{R}^n$. S is called convex if a straight-line segment joining any two points in S is contained entirely in S.

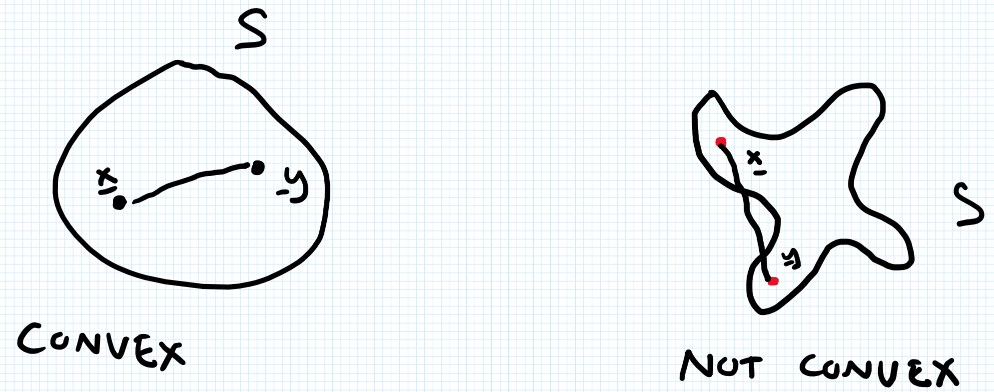

S is convex if, given any $\underline{x}$ and $\underline{y}$ in S.

S is convex if, given any $\underline{x}$ and $\underline{y}$ in S, the line segment

$$\underline{x}(t) = \underline{x}\,t + \underline{y}\,(1-t), \quad t \in [0,1]$$

is contained entirely in S.

Examples of convex sets — Unit Ball

Dot-product notation:

$$\langle \underline{x}, \underline{y} \rangle = \sum_{i=1}^{n} x_i y_i$$

$$\|\underline{x}\|_2^2 = \langle \underline{x}, \underline{x} \rangle = \sum_{i=1}^{n} x_i^2$$

Unit ball:

$$B = \{ \underline{x} \in \mathbb{R}^n \mid \|\underline{x}\|_2^2 \le 1 \}.$$

Theorem: B is a convex set.

Proof: Let $\underline{x}$ and $\underline{y} \in B$.

Form the line segment

$$\underline{x}(t) = \underline{x}\,t + \underline{y}\,(1-t), \quad t \in [0,1]$$

To show: $\|\underline{x}(t)\|_2^2 \le 1 \quad \forall \; t \in [0,1]$.

By direct computation:

$$\|\underline{x}(t)\|_2^2 = \langle \underline{x}(t), \underline{x}(t) \rangle$$

$$= \langle \underline{x}\,t + \underline{y}\,(1-t), \underline{x}\,t + \underline{y}\,(1-t) \rangle$$

$$= t^2 \langle \underline{x}, \underline{x} \rangle + t(1-t)\langle \underline{x}, \underline{y} \rangle$$

$$+ \ t(1-t) \langle \underline{y}, \underline{x} \rangle$$
$$+ \ (1-t)^2 \langle \underline{y}, \underline{y} \rangle$$

$$= t^2 \underbrace{\|\underline{x}\|_2^2}_{\leq 1} + 2t(1-t)\langle \underline{x}, \underline{y} \rangle + (1-t)^2 \underbrace{\|\underline{y}\|_2^2}_{\leq 1}$$

$$\leq t^2 \cdot 1 + 2t(1-t)\underbrace{\langle \underline{x}, \underline{y} \rangle}_{} + (1-t)^2 \cdot 1$$

C.S.

$$\leq t^2 + 2t(1-t)\|\underline{x}\|_2 \|\underline{y}\|_2 + (1-t)^2$$

$$\leq t^2 + 2t(1-t) + (1-t)^2$$

$$= \cancel{t^2} + \cancel{2t} - \cancel{2t^2} + 1 - \cancel{2t} + \cancel{t^2}$$

$$= 1.$$

Hence: $\|\underline{x}(t)\|_2^2 \leq 1 \quad \forall \ t \in [0,1]$.

Hence, $\underline{x}(t) \in B \quad \forall \ t \in [0,1]$

Hence, B is convex.

Polyhedra :

$$S = \{ \underline{x} \in \mathbb{R}^n \mid A\underline{x} = \underline{b}, \quad C\underline{x} \leq \underline{d} \}$$

where A and C are matrices of appropriate dimension and $\underline{b}$ and $\underline{d}$ are vectors.

$C\underline{x} \leq \underline{d}$ means $[C\underline{x}]_i \leq d_i$

S is a convex set (Homework).

Convex Function :

Let $f: S \to \mathbb{R}$, where $S \subset \mathbb{R}^n$.

The function f is convex if:

- S is a convex set
- The following relation holds:

$$f(t\underline{x} + (1-t)\underline{y}) \leq t f(\underline{x}) + (1-t) f(\underline{y}) \quad \forall t \in [0,1],$$

for all $\underline{x}, \underline{y} \in S$.

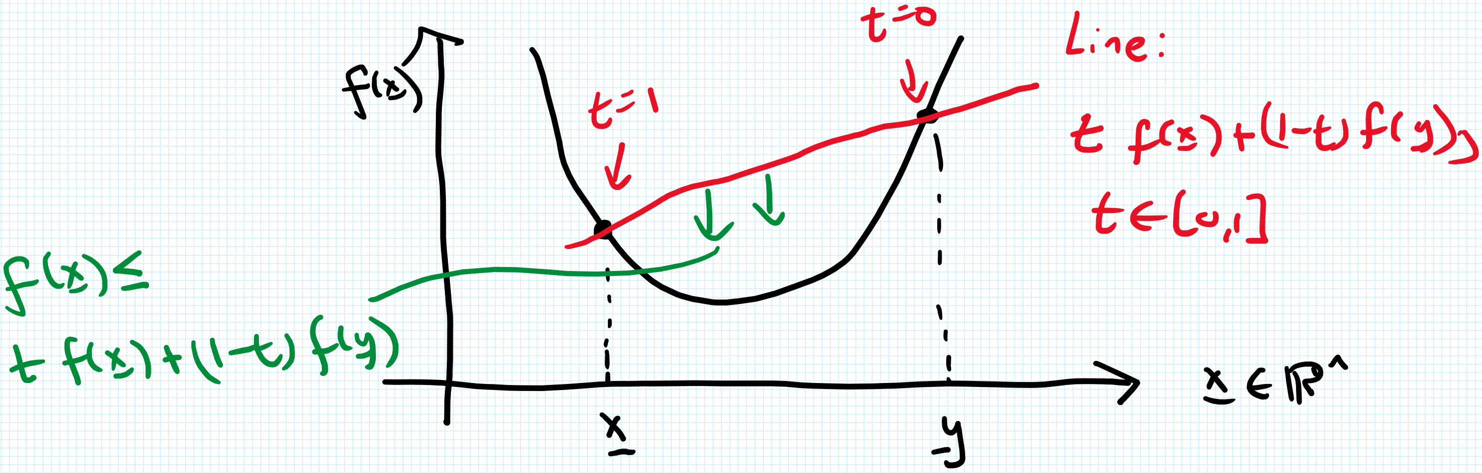

Some convex functions:

- Let $S = \mathbb{R}^n$. Then the linear function

$$f(\underline{x}) = \langle \underline{b}, \underline{x} \rangle + \alpha$$

is convex. Here, $\underline{b}$ is a constant vector and α is a constant scalar.

- Let $S = \mathbb{R}^n$. Then the following function (quadratic function) is convex:

$$f(\underline{x}) = \langle \underline{x}, H\underline{x} \rangle \qquad (1)$$

Where H is a symmetric semi-definite matrix.

Theorem: The function in (1) is convex.

Proof: Let $\underline{x}$ and $\underline{y}$ be two points in $\mathbb{R}^n$. Form the line segment $\underline{x}(t) = t \cdot \underline{x} + (1-t) \cdot \underline{y}$.

Look at

$$f(t\underline{x} + (1-t)\cdot\underline{y}) =$$

$$\langle t\underline{x} + (1-t)\underline{y}, H(t\underline{x} + (1-t)\underline{y}) \rangle$$

$$= t^2 \langle \underline{x}, H\underline{x} \rangle + t(1-t)\langle \underline{x}, H\underline{y} \rangle$$
$$+ t(1-t)\langle \underline{y}, H\underline{x} \rangle$$
$$+ (1-t)^2 \langle \underline{y}, H\underline{y} \rangle$$

Look at

$$\langle \underline{y}, H\underline{x} \rangle = \langle H\underline{x}, \underline{y} \rangle$$
$$= \langle \underline{x}, H^T\underline{y} \rangle$$
$$= \langle \underline{x}, H\underline{y} \rangle \quad \ldots H \text{ symmetric}$$

Hence:

$$f(t\underline{x} + (1-t)\underline{y}) =$$
$$t^2 \langle \underline{x}, H\underline{y} \rangle + 2t(1-t)\langle \underline{x}, H\underline{y} \rangle + (1-t)^2 \langle \underline{y}, H\underline{y} \rangle \quad (*)$$

Look at

$$t f(\underline{x}) + (1-t) f(\underline{y})$$
$$= t\langle \underline{x}, H\underline{x} \rangle + (1-t)\langle \underline{y}, H\underline{y} \rangle \quad (**)$$

Take $(*) - (**)$:

$$t^2 \langle \underline{x}, H\underline{x} \rangle + 2t(1-t)\langle \underline{x}, H\underline{y} \rangle + (1-t)^2 \langle \underline{y}, H\underline{y} \rangle$$

$$t^2\langle \underline{x}, H\underline{x}\rangle + 2t(1-t)\langle \underline{x}, H\underline{y}\rangle + (1-t)^2\langle \underline{y}, H\underline{y}\rangle$$
$$- t\langle \underline{x}, H\underline{x}\rangle \qquad\qquad - (1-t)\langle \underline{y}, H\underline{y}\rangle$$

$$= (t^2-t)\langle \underline{x}, H\underline{x}\rangle + 2t(1-t)\langle \underline{x}, H\underline{y}\rangle$$
$$+ \left[(1-t)^2 - (1-t)\right]\langle \underline{y}, H\underline{y}\rangle$$

$$= -t(1-t)\langle \underline{x}, H\underline{x}\rangle + 2t(1-t)\langle \underline{x}, H\underline{y}\rangle$$
$$+ (1-t)\left[\cancel{1} - t - \cancel{1}\right]\langle \underline{y}, H\underline{y}\rangle$$

$$= -t(1-t)\langle \underline{x}, H\underline{x}\rangle + 2t(1-t)\langle \underline{x}, H\underline{y}\rangle$$
$$-t(1-t)\langle \underline{y}, H\underline{y}\rangle$$

$$= -t(1-t)\left[\langle \underline{x}, H\underline{x}\rangle - 2\langle \underline{x}, H\underline{y}\rangle + \langle \underline{y}, H\underline{y}\rangle\right]$$

$$= \underbrace{-\underbrace{t(1-t)}_{\text{pos or zero}}\left[\underbrace{\langle (\underline{x}-\underline{y}), H(\underline{x}-\underline{y})\rangle}_{\text{pos or zero}}\right]}_{\text{neg or zero}}$$

$$\Rightarrow (*) - (**) \leq 0.$$

$$\Rightarrow (*) \leq (**)$$

$$\Rightarrow f(t\underline{x} + (1-t)\underline{y}) \leq tf(\underline{x}) + (1-t)f(\underline{y}) \ \forall t \in [0,1]$$

$\Rightarrow f'\left(t\underline{x} + (1-t)\underline{y}\right) \leq t f(\underline{x}) + (1-t) f(\underline{y}) \; \forall t \in [0,1]$

Hence, f is convex. ∎

Last definitions:

OP: $\min_{\underline{x} \in S} f(\underline{x})$, $S \subset \mathbb{R}^n$

A solution is written as $\underline{x}_*$.

$$\underline{x}_* = \arg\min_{\underline{x} \in S} f(\underline{x}).$$

Definition: $\underline{x}_*$ is a global minimizer if:

$$f(\underline{x}_*) \leq f(\underline{y}) \text{ for all } \underline{y} \in S.$$

$\underline{x}_*$ is a local minimizer if there exists a neighbourhood $N \subset S$ such that $\underline{x}_* \in N$ and such that $f(\underline{x}_*) \leq f(\underline{y}) \; \forall \underline{y} \in N$.

$\underline{x}_*$ is a strict local minimizer if there exists a neighbourhood $N \subset S$ such $\underline{x}_* \in N$ and such that

$$f(\underline{x}_*) < f(\underline{y}) \quad \forall \; \underline{y} \neq \underline{x}_* \in N.$$

$\underline{x}_*$ is an isolated local minimizer if there exists a a neighbourhood N of $\underline{x}_*$ such that $\underline{x}_*$ is the only minimum in N.

Pathological example:

$$f(x) = \begin{cases} x^4 \cos(\frac{1}{x}) + 2x^4, & x \neq 0 \\ 0 & x = 0 \end{cases}$$

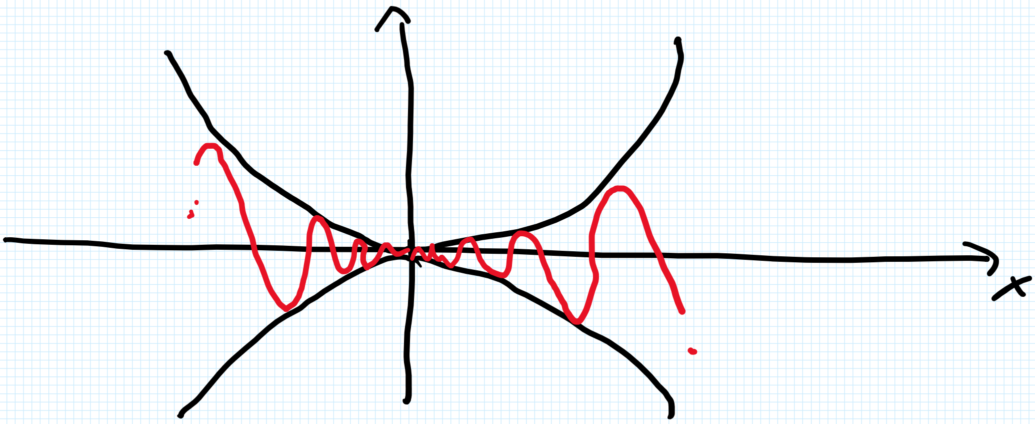

$x_* = 0$ is a minimizer but it is not isolated: you can always find another minimizer arbitrarily close to zero.

Necessary conditions for optimality (f differentiable) (§2.2)

Rigorously in 1D ($n=1$).

Taylor's Remainder Theorem: Let $f: \mathbb{R} \to \mathbb{R}$ be continuous with continuous first derivative.

Take $h > 0$. Then there exists $\eta \in (0,h)$ such that

$$f(x+h) = f(x) + f'(x+\eta) \cdot h$$

If f is twice continuously differentiable, then ther exists $\eta \in (0,h)$ such that:

$$f(x+h) = f(x) + f'(x)h + \frac{1}{2} f''(x+\eta) h^2.$$

Theorem 2.2 (First-order necessary condition).

Let x_* solve the OP. Then $f'(x_*) = 0$.

Let x_* solve the OP. Then $f'(x_*) = 0$.

Proof by contradiction: Assume for contradiction that $f'(x_*) < 0$. By continuity, there exists a $\delta > 0$ such that

$$f'(x) < 0 \quad , \quad x \in (x_* - \delta, x_* + \delta)$$

Look at $x = x_* + h$, where $|h| < \delta$.

Apply Taylor's Remainder Theorem:

$$f(x_* + h) = f(x_*) + f'(x_* + \eta) h, \quad \text{where}$$

$$|\eta| < |h| < \delta$$

$x_* + \eta$ is in the interval $(x_* - \delta, x_* + \delta)$ where $f'(x)$ is negative. Thus, take:

$$h = -\kappa f'(x_* + \eta) \quad , \quad \kappa > 0.$$

Hence:

$$f(x) = f(x_*) - \kappa |f'(x_* + \eta)|^2$$

Hence,

$$f(x) < f(x_*)$$ Contradiction, since x_* is the minimizer.

The calculation for $f'(x_*) > 0$ is the same – a contradiction arises in both cases, so we must have:

$$f'(x_*) = 0.$$

The result extends to n dimensions:

If $\underline{x}_*$ is a minimizer, then $\nabla f(\underline{x}_*) = 0$.

Theorem (Second-order necessary condition). Let x_* solve the OP $(n=1)$. Then $f''(x_*) \geq 0$.

Proof: By contradiction. Assume that $f''(x_*) < 0$.

By continuity, there exists a $\delta > 0$ such that

$$f''(x) < 0 \quad \forall x \in I = (x_* - \delta, x_* + \delta)$$

Consider $x = x_* + h$, where $|h| < \delta$, such that $x \in I$.

$$f(\overset{x_*+h}{x}) = f(x_*) + \underbrace{\cancel{f'(x_*)}}^{x_* \text{ is min}} h + \tfrac{1}{2} f''(x_* + \eta) h^2,$$

where $|\eta| < |h| < \delta$. Hence, $x_* + \eta \in I$.

$$\Rightarrow f(x) = f(x_*) + \frac{1}{2} \underbrace{f''(x_* + \eta)}_{<0} h^2$$

$\Rightarrow f(x) < f(x_*)$ Contradiction, since x_* is the minimizer.

Thus, to avoid the contradiction, we must have:

$$f''(x_*) \geq 0 \qquad \blacksquare$$

Extension to n dimensions: If x_* is a minimizer, then:

- $\nabla f(x_*) = 0$

- $H_{ij} = \left(\frac{\partial^2 f}{\partial x_i \partial x_j} \right)_{x_*}$ is positive-semi-definite.

Recall: H is positive-semi-definite means:

$$\langle \xi, H\xi \rangle \geq 0 \quad \forall \xi \in \mathbb{R}^n.$$

H is positive-definite means:

$$\langle \xi, H\xi \rangle > 0 \quad \forall \xi \neq 0 \in \mathbb{R}^n.$$

Necessary conditions:

x_* is a minimizer $\Longrightarrow$ $Df(x_*) = 0$, H_{ij} positive-semi-definite.

Sufficient conditions

x_* is a minimizer $\Longleftarrow$??

Theorem 2.7: Suppose that f is twice differentiable with continuous second derivatives. If $Df(x_*) = 0$ and if the Hessian matrix H is positive-definite, then x_* is a strict local minimizer.

Remark: Insisting on H being strictly positive-definite rules out saddle points and degenerate critical points (e.g. $f(x) = x^4$).

For convex functions, the situation is much nicer:

EXAM

Theorem: When f is convex, any local minimizer is a global minimizer. Also, if f is differentiable, then any stationary point ($\nabla f = 0$) is a global minimizer.