

LECTURE NOTES

INTRODUCTION TO OPTIMIZATION

FALL SEMESTER 2022

Roland Herzog*

2023-03-27

*Interdisciplinary Center for Scientific Computing, Heidelberg University, 69120 Heidelberg, Germany
(roland.herzog@iwr.uni-heidelberg.de, <https://scoop.iwr.uni-heidelberg.de/team/rherzog/>).

Material for 4 lectures of about 2 hours each.

Contents

§ 1	Introduction	5
§ 2	Unconstrained Optimization and Parameter Estimation	9
§ 2.1	Algorithmic Sketches for Generic Unconstrained Problems	9
§ 2.2	Algorithmic Sketches for Unconstrained Least-Squares Problems	11
§ 3	Convex Optimization	17
§ 4	Constrained Optimization	23
§ 5	Solutions	30
§ 5.1	Solution of Example 2.2 (Rosenbrock) using CASADi (Problem-Based)	30
§ 5.2	Solution of Example 2.2 (Rosenbrock) using CASADi (Opti Interface-Based)	31
§ 5.3	Solution of Example 2.3 (Airplane) using CASADi (Problem-Based)	32
§ 5.4	Solution of Example 4.1 (Production) using MATLAB's Optimization Toolbox (Solver-Based)	33
§ 5.5	Solution of Example 4.1 (Production) using MATLAB's Optimization Toolbox (Problem-Based)	34
§ 5.6	Solution of Example 4.1 (Production) using <code>scipy.optimize</code>	35
§ 5.7	Solution of Example 4.1 (Production) using CASADi (Problem-Based)	36
§ 5.8	Solution of Example 4.1 (Production) using CASADi (Opti Interface-Based)	37
§ 5.9	Solution of Example 4.12 (Post-Office) using CASADi (Opti Interface-Based)	38

§ 1 INTRODUCTION

Optimization is about solving problems of the form

$$\left. \begin{array}{ll} \text{Minimize} & f(x) \quad \text{where } x \in \mathbb{R}^n \quad \text{(objective (function))} \\ \text{subject to} & g_i(x) \leq 0 \quad \text{with } i = 1, \dots, n_{\text{ineq}} \quad \text{(inequality constraints)} \\ & \text{and } h_j(x) = 0 \quad \text{with } j = 1, \dots, n_{\text{eq}}. \quad \text{(equality constraints)} \end{array} \right\} \quad (1.1)$$

x is termed the **optimization variable**, **decision variable**, or simply the **variable** of the problem. Moreover, we consider problems where

- the functions $f, g_i, h_j: \mathbb{R}^n \rightarrow \mathbb{R}$ are sufficiently smooth,
- the number n_{ineq} of inequality constraints and the number n_{eq} of equality constraints are finite (possibly zero).

More precisely, problems of type (1.1) fall into the realm of continuous **optimization**. In case some or all components of x are to be sought in a discrete set (such as the integers \mathbb{Z}), we have a mixed discrete-continuous problem or even a fully discrete problem, which we do not consider in this course.

Solving an optimization problem usually includes the following steps:

(1) Modeling the problem.

This step can be one of the most complex ones in the entire process. In particular, we need to answer the following questions:

- What do we use as our optimization variable(s)? What is the significance of each component of x ?
- How do we formulate the objective function? What is its significance?
- Do we need to formulate constraints? How do we model them? What is the significance of each of the constraints?
- Are there any constants in the problem? What are their values and what is their significance?

(2) Understanding the problem's characteristics.

There are various dimensions along which you can classify optimization problems. Here are some examples:

- What is the dimension n of the optimization variable $x \in \mathbb{R}^n$?
- Is the problem constrained or unconstrained?
- Is the objective linear or nonlinear? Are the constraints (if any) linear or nonlinear?
- Is the problem convex or non-convex?

The answer to these questions has an impact on the next step, the selection of a suitable solver. The more specifically the solver fits the problem, the more efficient the solution process usually

is. For instance, a linear optimization problem (where the objective and all constraints are affine functions of x) is best solved using a solver tailored to linear optimization problems (such as an implementation of the simplex method). It is “overkill” to solve a linear optimization problem using a solver capable of solving general nonlinear constrained problems and the more general optimization procedure will typically perform poorer than specialized routines (although it should usually succeed).

(3) Picking a numerical solver.

Pick your solver depending on the problem’s characteristic as well as availability/license and your preferred programming language.

The following resources can help you identify suitable solvers:

- Hans Mittelmann’s [Decision Tree for Optimization Software](#)
- the [Optimization Tree](#) of the NEOS project (a free internet-based service for solving optimization problems)
- the list of available solvers in the [COIN-OR project](#)
- the respective documentation of various optimization libraries, such as
 - MATLAB’s Optimization Toolbox
 - the [scipy.optimize](#) PYTHON package
 - the [CASADI](#) package, which has MATLAB, PYTHON and C++ bindings

(4) Implementing the functions describing the problem for the solver.

You will need to implement the objective and constraint functions in a format suitable for the solver selected in the respective programming language (such as MATLAB, PYTHON, JULIA, or C++ for instance).

There are also specific [modeling languages](#) for optimization problems, e. g., AMPL, which are, in principle, independent of the solver. However, not all solvers will accept all input formats.

Finally, some solver packages such as CASADI (PYTHON, MATLAB) and JUMP (JULIA), include their own modeling paradigms which facilitate the formulation of optimization problems.

(5) Running the solver.

Launch the solver selected for the problem. In some cases, you may want to modify the default settings (such as subalgorithms, tolerances, maximum number of iterations etc.). Most solvers will accept an initial guess (often termed x_0) for the solution, which is often beneficial to set to a reasonable value.

(6) Checking and interpreting the result.

Carefully check the flag or message returned by the solver, i. e., whether it has converged or not. When successful, interpret and verify the result. Check that the solution “makes sense”. Are all the constraints satisfied as expected? Are you happy with the result, or did you perhaps expect something different? Did you forget to include certain constraints? (A classical example is that you expect the solution to be non-negative but forget to model this constraint.)

The above steps are not necessarily run through sequentially. For instance, it might turn out later in the process that a different formulation of the problem than originally selected is more favorable for the solver, so one needs to go back. Or we might figure out from the solution that we forgot to model certain constraints, which we then include and run the solver again.

It is also important to understand that, given a problem in text form, the above steps do not have a predetermined outcome. It is up to us to **model the problem** so that we can identify (or design ourselves) a numerical solver that can efficiently solve the problem. Your modeling skills will get better with practice, and this process is partially based on trial-and-error.

Example 1.1 (Choice of optimization variables). *This examples illustrates that there is no single correct answer to modeling a problem. For instance, what we choose to be our optimization variables is a modeling decision. Consider, for instance, an investor, who seeks to distribute a certain amount of money M between a bank account with fixed interest and a stock investment with uncertain future stock prices. The investor's objective may be to maximize the expected amount of their wealth at a certain time in the future.*

The investor may choose optimization variables

$$\begin{aligned} x_1 & \text{ amount of money invested in the bank account,} \\ x_2 & \text{ amount of money invested in the stock.} \end{aligned}$$

We are not specifying the problem any further here, but we mention that it is likely going to contain the constraints $x_1 \geq 0$, $x_2 \geq 0$ and $x_1 + x_2 = M$.

Alternatively, instead of determining how much money to put in either investment, the investor may describe their optimization variable using a single decision, namely what fraction ("percentage") of the money M is going to be invested in the bank account. That is, the investor may choose their optimization variable as

$$x_1 \text{ fraction of } M \text{ to be invested in the bank account,}$$

which is subject to the constraint $0 \leq x_1 \leq 1$. Naturally, the remainder fraction $1 - x_1$ of M is going to be invested in the stock.

We now discuss what it means to solve an optimization problem, i. e., to find and recognize minimizers.

Definition 1.2 (Basic concepts).

(i) *The set of points $x \in \mathbb{R}^n$ satisfying all the constraints, i. e.,*

$$F := \{x \in \mathbb{R}^n \mid g_i(x) \leq 0 \text{ for all } i = 1, \dots, n_{\text{ineq}} \text{ and } h_j(x) = 0 \text{ for all } j = 1, \dots, n_{\text{eq}}\} \quad (1.2)$$

*is termed the **feasible set**. Every $x \in F$ is called a **feasible point**.*

(ii) *The inequality $g_i(x) \leq 0$ is said to be **active** at the point x in case of $g_i(x) = 0$. It is called **inactive** in case $g_i(x) < 0$. It is called **violated** in case $g_i(x) > 0$.*

(iii) A point $x^* \in F$ is a **global minimizer** or **global solution** of (1.1) if

$$f(x^*) \leq f(x) \text{ for all } x \in F.$$

(iv) A point $x^* \in F$ is a **local minimizer** or **local solution** of (1.1) if there exists a (potentially small) neighborhood of x^* ¹ $U(x^*)$ such that

$$f(x^*) \leq f(x) \text{ for all } x \in F \cap U(x^*).$$

While Definition 1.2 states what minimizers are, the definition is not very helpful when it comes to checking whether a given point x^* is indeed, say, a local minimizer. In order to verify this on the grounds of the definition would require us to compare the value of the objective $f(x^*)$ with the value of the objective $f(x)$ at an (uncountable!) number of nearby points x ! This is clearly impractical.

Therefore, we need to resort to different ways of verifying that a certain point is a minimizer, and also of excluding certain points from the list of candidate minimizers. To this end, we can use derivative-based optimality conditions. Optimality conditions come in two flavors:

- (1) A **necessary optimality condition** (NC) is a statement of the following form: “Suppose that x^* is a local minimizer of problem (1.1). Then (NC) holds at x^* .”

Necessary optimality conditions can be used in the following ways:

- (a) Finding points $x^* \in \mathbb{R}^n$ that satisfy (NC) means to compile a list of candidates for local minimizers. Those can then be further examined by other means.
- (b) A necessary optimality condition can also be read as:² “Suppose that (NC) does *not* hold at x^* , then x^* cannot be a local minimizer.” Consequently, necessary optimality conditions can be used to exclude points from the search for local minimizers.
- (2) A **sufficient optimality condition** (SC) is a statement of the following form: “Suppose that (SC) holds at x^* . Then x^* is a local minimizer.”

Sufficient optimality conditions thus allow us to verify candidates as local minimizers.

We will discuss appropriate forms of necessary and sufficient optimality conditions for the particular type of problem in each of the following sections.

¹Note that the definition is equivalent if we were to replace “neighborhood of x^* ” by “ball around x^* ”, i. e., $B_\varepsilon(x^*) = \{x \in \mathbb{R}^n \mid \|x - x^*\|_2 < \varepsilon\}$.

²This is called the contraposition of the original statement.

§ 2 UNCONSTRAINED OPTIMIZATION AND PARAMETER ESTIMATION

As the name suggests, unconstrained optimization is about minimizing an objective in the absence of any equality or inequality constraints. In this case, our generic problem (1.1) reduces to

$$\text{Minimize } f(x) \quad \text{where } x \in \mathbb{R}^n. \quad (2.1)$$

For this problem, “derivative equals zero” is a necessary optimality condition. Points which satisfy $f'(x) = 0$ are called **stationary points**.

Theorem 2.1 (Necessary optimality condition for unconstrained problems).

Suppose that x^ is a local minimizer for (1.1). Then x^* is a stationary point, i. e., $f'(x^*) = 0$ holds.*

Sufficient optimality conditions are usually based on second-order derivatives, and we do not state them here. Instead, we proceed to sketch some algorithms. Indeed, algorithms usually find stationary points, which may or may not be local minimizers. Nevertheless, we use the term “optimization algorithms” and continue to speak of “minimizing the objective”.

§ 2.1 ALGORITHMIC SKETCHES FOR GENERIC UNCONSTRAINED PROBLEMS

Optimization algorithms proceed in a sequential way and break down problem (1.1) into a sequence of simpler problems. In doing so, they generate a sequence $x^{(k)}$ of iterates, which (under some conditions) converges to a stationary point x^* .

The vast majority of algorithms for (1.1) proceeds as follows. At the current iterate $x^{(k)}$, they form a **quadratic model** of the objective. A quadratic model is a quadratic (i. e., second-order) polynomial that shares some properties with the true objective in the vicinity of the current iterate $x^{(k)}$. The quadratic model at $x^{(k)}$ is of the form (see Figure 2.1 for an illustration in 1D)

$$q^{(k)}(x) := \underbrace{f(x^{(k)})}_{\text{constant term}} + \overbrace{f'(x^{(k)})(x - x^{(k)})}^{\text{linear term}} + \underbrace{\frac{1}{2}(x - x^{(k)})^\top H^{(k)}(x - x^{(k)})}_{\text{quadratic term}}. \quad (2.2)$$

The matrix $H^{(k)}$ is called the **model Hessian** and it is designed to be symmetric.

At $x = x^{(k)}$, the model $q^{(k)}$ shares its function value and its derivative with the true objective f . (**Quiz 2.1:** Can you show this?)

Algorithms proceed by minimizing the current quadratic model $q^{(k)}$, or at least they look for a stationary point of the model:

$$[q^{(k)}]'(x) = f'(x^{(k)}) + (x - x^{(k)})^\top H^{(k)} \stackrel{!}{=} 0.$$

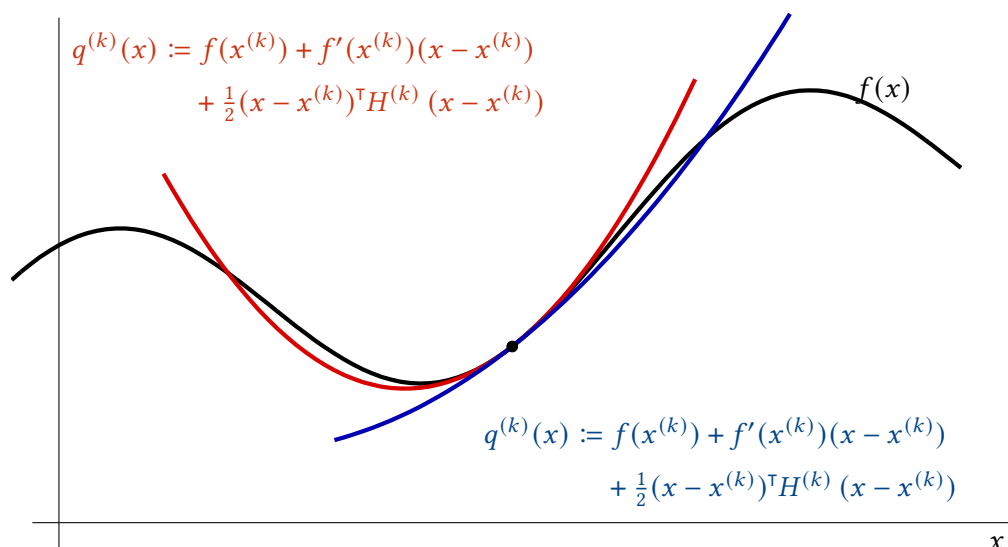


Figure 2.1: A function $f := \mathbb{R} \rightarrow \mathbb{R}$ and two possible quadratic models $q^{(k)}$ about the point $x^{(k)}$ (black dot). The red model is the second-order Taylor polynomial of f at $x^{(k)}$. In other words, the red model uses $H^{(k)} = f''(x^{(k)})$. The blue model instead uses an “incorrect” Hessian $H^{(k)} \neq f''(x^{(k)})$.

This amounts to solving the linear(!) system of equations

$$H^{(k)}(x - x^{(k)}) = -\nabla f(x^{(k)}).$$

Usually this is written in the form

$$H^{(k)}d = -\nabla f(x^{(k)}) \quad (2.3)$$

whose solution gives us an update direction d in the space of optimization variables.

In practical algorithms, this direction is then scaled with a suitable step length $\alpha^{(k)} > 0$, which accounts for the deviation of the quadratic model from the original objective and gives us the update

$$x^{(k+1)} := x^{(k)} + \alpha^{(k)} d.$$

The step length $\alpha^{(k)}$ has the purpose of making the step between consecutive iterates small when the model does not agree well with the true objective. This is done in an effort to obtain robust convergence properties regardless of the quality of the initial guess provided by the user. Since $\alpha^{(k)}$ is determined by a line search procedure, in which trial step sizes are proposed and evaluated until a satisfactory one is found, one speaks of a **line search algorithm**. An alternative concept is to limit the size of d in the first place by constraining the norm $\|d\|_2 \leq \Delta^{(k)}$ when minimizing the model (2.2). The solution of this problem is more complex than solving a linear system (2.3). This leads to **trust-region algorithms**, and $\Delta^{(k)}$ is the trust-region radius.

Depending on the choice of the model Hessians $H^{(k)}$ implemented in an algorithm, we refer to it by different names:³

³Combined with the choice above, one then obtains a *line-search quasi-Newton method*, or a *trust-region Newton algorithm*, for instance. The combination *trust-region gradient method* is, however, uncommon.

gradient method (steepest descent)	$H^{(k)} = \text{Id}$ (identity matrix)	linear convergence
preconditioned gradient method	$H^{(k)} = M$	linear convergence
quasi-Newton method	$H^{(k)}$ varies	often superlinear convergence
Newton method	$H^{(k)} = f''(x^{(k)})$	superlinear convergence

Gradient methods are simple in the sense that the linear system (2.3) solved in each iteration has a constant coefficient matrix, which can be exploited. Gradient methods, however, typically exhibit slow convergence. On the other hand, Newton methods converge faster once they are near a minimizer; however, the evaluation of the objective's second derivative matrix in every iteration is usually costly. In practice, quasi-Newton methods are a good compromise, in which $H^{(k)}$ tries to capture some of the properties of the objective's second derivative without actually computing it.

We close this section by mentioning what functions a user has to implement in order to run an algorithm for solving an unconstrained problem (2.1):

objective	$f(x)$	
derivative	$f'(x)$	optional (finite difference fallback)
2nd derivative	$f''(x)$ or $f''(x) d$	optional, only for Newton methods

The evaluation of the first-order derivative $f'(x)$, although required by all algorithms, is usually optional for the user. When it is not provided, the algorithm will fall back to an internal approximation of $f'(x)$ by finite differences. While this may sound convenient, it may be very time consuming and introduce inaccuracies. If possible, the user should provide first-order derivatives. To facilitate the task, they may use **algorithmic differentiation** (AD) for this purpose. This is beyond the scope of these notes.

Example 2.2 (Rosenbrock problem). The **Rosenbrock problem** is a classical test example in unconstrained optimization:

$$\text{Minimize } f(x) := (a - x)^2 + b(y - x^2)^2 \quad \text{where } (x, y) \in \mathbb{R} \times \mathbb{R}$$

It has a global minimizer at $(x, y) = (a, a^2)$, with an optimal function value of 0. A typical choice of the parameters is $(a, b) = (1, 100)$. A plot of the function can be found in Figure 2.2.

§ 2.2 ALGORITHMIC SKETCHES FOR UNCONSTRAINED LEAST-SQUARES PROBLEMS

Many unconstrained optimization problems are in fact of the form

$$\text{Minimize } f(x) := \sum_{i=1}^M [r_i(x)]^2 = \|r(x)\|_2^2 \quad \text{where } x \in \mathbb{R}^n. \quad (2.4)$$

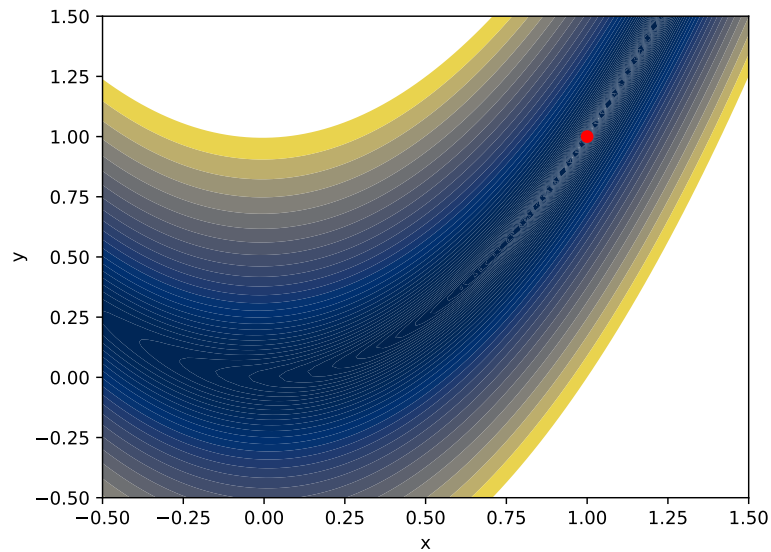


Figure 2.2: Contour plot of the Rosenbrock “banana” function (Example 2.2) with parameters $(a, b) = (1, 100)$, and its global minimizer.

These are called **least-squares problems**. The primary source of least-squares problems lies in **parameter estimation**, also known as **parameter identification**, **model calibration**, **curve fitting**, or **data assimilation**.

In parameter estimation one is given a model function m which describes a relation $\eta = m(\xi)$ between *independent variables* (model input) $\xi \in \mathbb{R}^r$ and *dependent variables* $\eta \in \mathbb{R}$ (model output). The model in fact contains yet unknown parameters $x \in \mathbb{R}^n$, i. e., we have a relation of the form

$$\eta = m(x; \xi).$$

Each measurement pair (ξ_i, η_i) contributes a **residual**

$$r_i(x) := \underbrace{m(x; \xi_i)}_{\text{model prediction}} - \underbrace{\widehat{\eta}_i}_{\text{measurement}}, \quad (2.5)$$

which describes the discrepancy between the value predicted by the model and the actual measurement η_i pertaining to the i -th input ξ_i .

The least-squares problem (2.4) describes our desire to choose the parameters $x \in \mathbb{R}^n$ in such a way that the model will fit the available measurement pairs $(\xi_i, \eta_i) \in \mathbb{R}^r \times \mathbb{R}$, $i = 1, \dots, M$, in the best possible way, in the sense of least squares.

There are statistical reasons why – at least for measurements η_i perturbed by additive random measurement errors, which are uncorrelated and identically normally distributed – the solution of the least-squares problem (2.1) yields the best estimate of the parameters. When, more generally, the

measurement errors are still uncorrelated but may have non-identical variances σ_i^2 , we should consider the **weighted least-squares problem**

$$\text{Minimize } f(x) := \sum_{i=1}^M \frac{1}{\sigma_i^2} [r_i(x)]^2 = \|r(x)\|_{\Sigma^{-1}}^2 \quad \text{where } x \in \mathbb{R}^n, \quad (2.6)$$

where $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_M^2)$ is the measurement covariance matrix. However, we do not specifically consider the form (2.6) here. After all, (2.6) can be brought into the form (2.4) simply by scaling the residuals by $1/\sigma_i$.

We could use any of the algorithmic ideas from § 2.1 to solve (2.4). However, it is usually more efficient to use a different algorithm that exploits the particular structure of the objective in (2.4). In order to understand this, we determine the first- and second-order derivatives of the objective from (2.4) using the chain rule. We expect that the derivative of the residual function r will appear, i. e., its Jacobian, which we denote by r' or by J :

$$J(x) := \left[\frac{\partial}{\partial x_j} r_i(x) \right]_{ij} = \begin{bmatrix} | & & | \\ \frac{\partial}{\partial x_1} r(x) & \cdots & \frac{\partial}{\partial x_n} r(x) \\ | & & | \end{bmatrix} = \begin{bmatrix} -\nabla r_1(x)^\top - \\ \vdots \\ -\nabla r_m(x)^\top - \end{bmatrix} \in \mathbb{R}^{M \times n}. \quad (2.7)$$

For cosmetic reasons it is practical to divide the objective in (2.4) by 2. This does, of course, not influence where the minimizers lie.

$$f(x) = \frac{1}{2} \|r(x)\|_2^2 = \frac{1}{2} r(x)^\top r(x) \quad \text{objective} \quad (2.8a)$$

$$\nabla f(x) = J(x)^\top r(x) = \sum_{i=1}^M r_i(x) \nabla r_i(x) \quad \text{first-order derivative} \quad (2.8b)$$

$$f''(x) = J(x)^\top J(x) + \sum_{i=1}^M r_i(x) r_i''(x) \quad \text{second-order derivative.} \quad (2.8c)$$

With the help of this data one *could* set up a quadratic model of the objective f at the current iterate $x^{(k)}$, determine a search direction $d^{(k)}$ by minimizing it, then find a step length $\alpha^{(k)}$, take the step etc. The additional knowledge, however, that we are dealing with a least-squares problem can be exploited even better algorithmically. We describe this using the popular **Levenberg-Marquardt method**⁴ as an example.

The Levenberg-Marquardt method sets up a certain quadratic model about the point $x^{(k)}$:

$$\begin{aligned} q_{\text{LM}}^{(k)}(d) &:= f(x^{(k)}) + \nabla f(x^{(k)})^\top d + \frac{1}{2} d^\top H_{\text{LM}}^{(k)} d \\ &= \frac{1}{2} r(x^{(k)})^\top r(x^{(k)}) + r(x^{(k)})^\top J(x^{(k)}) d + \frac{1}{2} d^\top H_{\text{LM}}^{(k)} d \quad \text{by (2.8)}. \end{aligned} \quad (2.9)$$

The Levenberg-Marquardt method uses

$$H_{\text{LM}}^{(k)} := J(x^{(k)})^\top J(x^{(k)}) + \lambda^{(k)} \text{Id} \quad (2.10)$$

as the model Hessian. Here $\lambda^{(k)} > 0$ is a positive number. This has the following advantages:

⁴proposed by Levenberg (1944), rediscovered by Marquardt (1963)

- (i) One uses the “good” first part $J(x)^\top J(x)$ of the true Hessian (2.8c). This part is at least positive semidefinite and it can be computed without further effort since we need $J(x)$ anyway for the first-order derivative $\nabla f(x)$.
- (ii) The “bad” part $\sum_{i=1}^M r_i(x) r_i''(x)$ of the Hessian (2.8c) is omitted. Since this part depends on the second-order derivatives $r_i''(x)$, it is expensive to evaluate, and it can potentially destroy the positive definiteness. Moreover we hope that the residuals $r_i(x^*)$ at the solution x^* will be small (near zero), i. e., we hope that we will have a good fit, and this part of the Hessian can be disregarded.
- (iii) Instead one adds an auxiliary term $\lambda^{(k)} \text{Id}$. In view of $\lambda^{(k)} > 0$, positive definiteness of the model Hessian $H_{\text{LM}}^{(k)}$ is ensured, which implies that the linear problems that need to be solved in the algorithms substeps are uniquely solvable.

As with the methods of § 2.1, one is interested in the minimizer of the quadratic model (2.9) in the k -th iteration. This minimizer is unique and it can be calculated by solving the following linear system of equations, compare (2.3),

$$H_{\text{LM}}^{(k)} d^{(k)} = -\nabla f(x^{(k)}). \quad (2.11)$$

In more explicit terms, see (2.10) and (2.8b), we can write

$$[J(x^{(k)})^\top J(x^{(k)}) + \lambda^{(k)} \text{Id}] d^{(k)} = -J(x^{(k)})^\top r(x^{(k)}). \quad (2.12)$$

The size of the Levenberg-Marquardt parameter $\lambda^{(k)}$ implicitly determines the length of the step d . (**Quiz 2.2:** What happens in the extreme cases $\lambda^{(k)} \rightarrow 0$ and $\lambda^{(k)} \rightarrow \infty$, respectively?) $\lambda^{(k)}$ itself is controlled depending on how well the model (2.9) was able to predict the actual decrease in the values of the objective (2.4).

We mention what functions a user has to implement in order to run a dedicated least-squares algorithm for solving a problem of type (2.4):

residual	$r(x)$
Jacobian	$J(x)$ optional (finite difference fallback)

Often times, the user has a choice (through different interfaces to the algorithm) to specify the model function m and the set of measurement pairs $(\xi_i, \eta_i) \in \mathbb{R}^r \times \mathbb{R}$, $i = 1, \dots, M$, and have the method itself evaluate the residual function with components (2.5) and Jacobian

$$J(x) = r'(x) = \begin{bmatrix} \frac{\partial}{\partial x_1} m(x; \xi_1) & \cdots & \frac{\partial}{\partial x_n} m(x; \xi_1) \\ \vdots & & \vdots \\ \frac{\partial}{\partial x_1} m(x; \xi_M) & \cdots & \frac{\partial}{\partial x_n} m(x; \xi_M) \end{bmatrix} = \begin{bmatrix} -\nabla m(x; \xi_1)^\top \\ \vdots \\ -\nabla m(x; \xi_M)^\top \end{bmatrix}.$$

In this case, the user provides

model	$m(x; \xi)$	
model gradient	$\nabla m(x; \xi)$	optional (finite difference fallback)

and the set of measurement pairs (ξ_i, η_i) . Again, when derivatives are not provided, most algorithms will fall back to finite differencing.

Example 2.3 (Parameter estimation). *The position (x, y) of an airplane should be determined by radio bearing. To this end, the directions towards several radio beacons are measured from the airplane. The positions (X_i, Y_i) of the beacons are known and the directions are measured as angles α_i to the x-axis:*

beacon i	beacon's position (X_i, Y_i)	angle α_i
1	(8, 6)	38°
2	(-3, -3)	220°
3	(1, 0)	222°
4	(8, -3)	300°

The aim is to find the position (x, y) of the airplane for which the discrepancies between the predicted and measured angles are as small as possible in the least squares sense.

In this example, we have $M = 4$ individual measurements. We require a model which maps the independent variable $\xi = (X, Y) \in \mathbb{R}^2$ to the dependent variable $(\eta = \alpha \in \mathbb{R})$, as a function of the unknown airplane position (x, y) . A model which achieves this is given by

$$m(\underbrace{(x, y)}_{\text{unknown parameter}}; \underbrace{(X, Y)}_{\text{independent variable}}) = \text{atan2}(X - x, Y - y) \cdot \frac{180^\circ}{\pi},$$

which returns the angle of the vector $\begin{pmatrix} X-x \\ Y-y \end{pmatrix}$ against the x-axis in degrees. In fact, since `atan2` returns angles in the interval $(-\pi, \pi]$, while our measurements are in 0° to 360° , we have to normalize `atan2`'s output to $[0, 2\pi)$ beforehand; see the listing in § 5.3 for details.

The solution obtained with the code from § 5.3 is shown in Figure 2.4.

End of Class 1

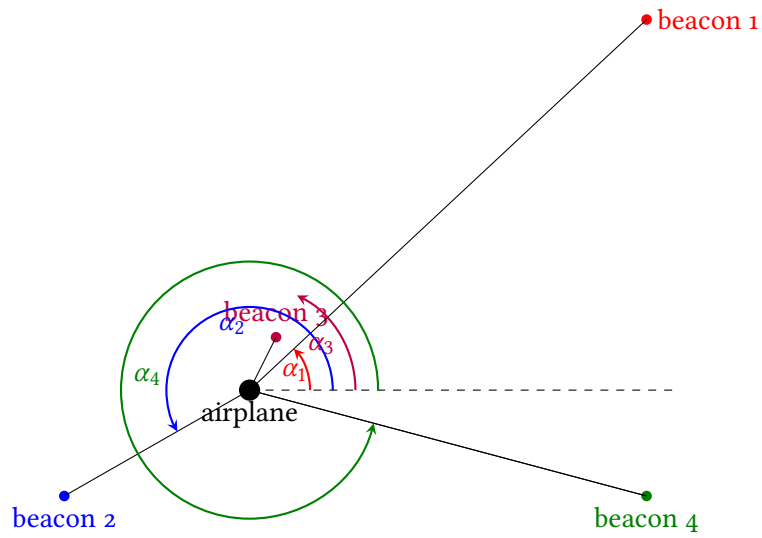


Figure 2.3: Illustration of the airplane position parameter estimation problem from [Example 2.3](#).

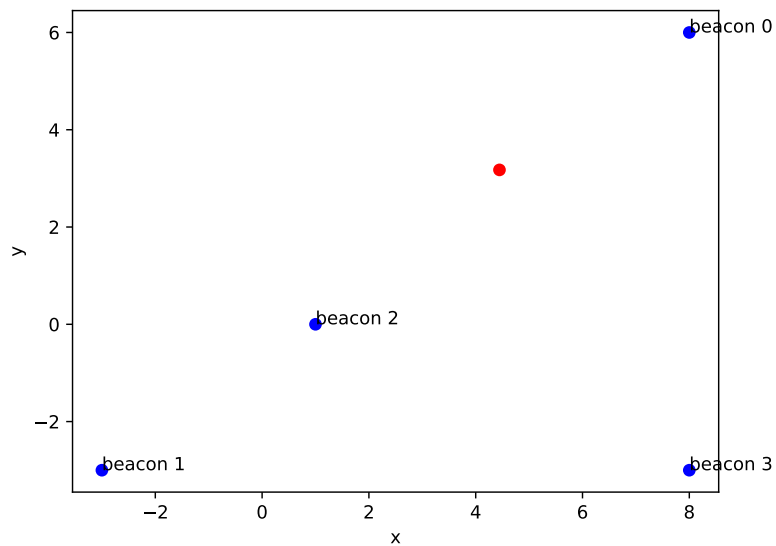


Figure 2.4: Solution of the airplane position parameter estimation problem from [Example 2.3](#).

§ 3 CONVEX OPTIMIZATION

In this section we consider another class of optimization problems:

$$\text{Minimize } f(x) \quad \text{where } x \in \mathbb{R}^n. \quad (3.1)$$

These problems feature a *convex* objective $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$. Due to the possibility of f taking the value ∞ , we can implicitly model constraints, since points with a function value $f(x) = \infty$ are not considered to be minimizers, since we exclude the pathological case $f \equiv \infty$.

Definition 3.1 (Convex function). *Suppose that $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$.*

(i) f is said to be **convex** if

$$f(\alpha x + (1 - \alpha) y) \leq \alpha f(x) + (1 - \alpha) f(y) \quad (3.2)$$

holds for all $x, y \in \mathbb{R}^n$ and all $\alpha \in [0, 1]$.

(ii) f is said to be **strictly convex** if

$$f(\alpha x + (1 - \alpha) y) < \alpha f(x) + (1 - \alpha) f(y) \quad (3.3)$$

holds for all $x \neq y \in \mathbb{R}^n$ and all $\alpha \in (0, 1)$.

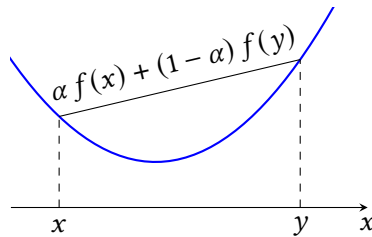


Figure 3.1: Example of a (strictly) convex function on \mathbb{R} .

The benefit of the objective being convex is clarified by the following theorem.

Theorem 3.2 (Fundamental theorem of convex optimization).

- (a) Every local minimizer of (3.1) is already a global minimizer.
- (b) The set of global minimizers of (3.1) is a convex set.⁵
- (c) When f is strictly convex, then (3.1) has at most one solution.

⁵A set $C \subseteq \mathbb{R}^n$ is said to be **convex** if the line joining any points in C remains in C , i. e.: for all $x, y \in C$ we have $\alpha x + (1 - \alpha) y \in C$ for all $\alpha \in [0, 1]$.

Consequently, in convex optimization one does not have to distinguish local from global minimizers.

We consider in this section optimization problems of a form more particular than the generic convex problem (3.1). Indeed, we consider composite problems of the form

$$\text{Minimize } f(x) + g(Ax) \quad \text{where } x \in \mathbb{R}^n. \quad (3.4)$$

In (3.4), $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ and $g: \mathbb{R}^m \rightarrow \mathbb{R} \cup \{\infty\}$ are both convex functions, and $A \in \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a matrix. Problems of type (3.4) are encountered in many practical problems, and they allow specialized algorithms of convex optimization to be used, compared to the generic convex problem (3.1).

Example 3.3 (Image denoising). *A classical problem in image denoising is as follows. Suppose we are given a noisy image (gray-scale for simplicity), which is represented as a matrix $Z \in \mathbb{R}^{n_1 \times n_2}$. The image height (number of pixels in vertical direction) is n_1 while n_2 is the width. We wish to remove the noise while retaining the image's content. In particular, we wish to preserve edges, i. e., sharp variations of the brightness values between neighboring pixels.*

A classical way to approach this task (Rudin, Osher, Fatemi, 1992) is to formulate it as an optimization problem:

$$\text{Minimize } \underbrace{\frac{1}{2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} (X_{i,j} - Z_{i,j})^2}_{\text{fidelity term}} + \beta \underbrace{\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} |X_{i,j} - X_{i,j+1}|}_{\text{horizontal jump}} + \beta \underbrace{\sum_{i=1}^{n_1-1} \sum_{j=1}^{n_2} |X_{i+1,j} - X_{i,j}|}_{\text{vertical jump}} \quad (3.5)$$

where $X \in \mathbb{R}^{n_1 \times n_2}$.

The first term ensures that the resulting image X does not deviate too much from the noisy image Z observed. The remaining terms penalize differences in brightness between all pairs of neighboring pixels, both in horizontal and in vertical directions. The parameter $\beta > 0$ balances both objectives.

Problem (3.5) can be written in the form (3.4) when we vectorize⁶ the matrix-valued optimization variable to become $x = \text{vec}(X) \in \mathbb{R}^{n_1 n_2}$. Likewise, we vectorize the observed image: $z = \text{vec}(Z)$. The first term in (3.5) then simply becomes $\frac{1}{2} \|x - z\|_2^2$. To account for the remaining terms, we require a matrix A which converts the pixel values x to the vector of horizontal differences, followed by the vector of vertical differences. Such a matrix A is of dimension $[n_1(n_2 - 1) + n_2(n_1 - 1)] \times [n_1 n_2]$. Each row contains precisely one entry +1 and one entry -1, all remaining entries are zero; see Figure 3.2. Finally, $g(y) = \beta \|y\|_1$ holds, i. e., $g(y)$ is β times the sum of the absolute values of all entries of the vector y .

The objective in Example 3.3 is convex since

$$f(x) := \frac{1}{2} \|x - z\|_2^2 \quad \text{and} \quad g(Ax) := \|Ax\|_1 \quad (3.6)$$

are both convex, and since sums and non-negative multiples of convex functions are again convex.

The algorithm we are considering exploits the additive and composite structure of (3.5). In order to derive it, we require the notion of the Fenchel conjugate of a function.

⁶Vectorization of a matrix simply means to stack the columns of the matrix on top of each other to form a long vector with all entries of the matrix.

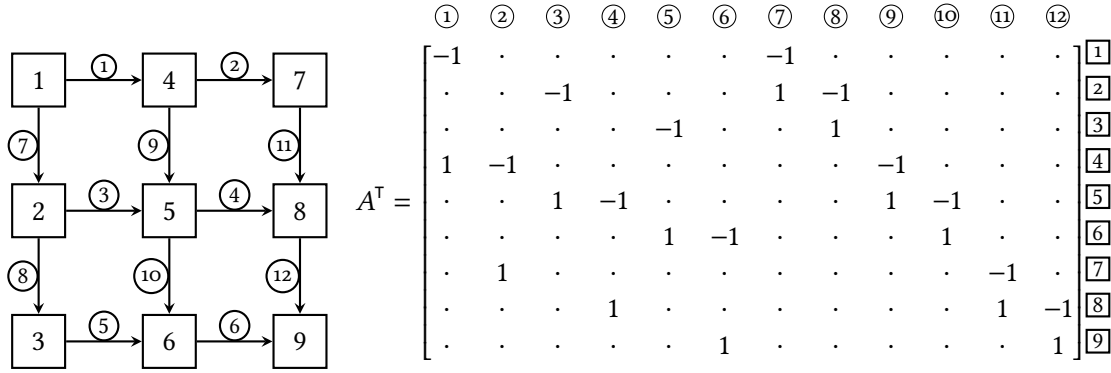


Figure 3.2: Illustration of the construction of the matrix A for an image denoising problem (Example 3.3). Its transpose A^T is the incidence matrix of the directed graph obtained from connecting neighboring pixels. For readability, zeros are written as \cdot .

Definition 3.4 (Fenchel conjugate). Suppose that $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ is a function. Then the **Fenchel conjugate** of f is $f^*: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\pm\infty\}$, defined by

$$f^*(\xi) := \sup_{x \in \mathbb{R}^n} \xi^T x - f(x). \quad (3.7)$$

The Fenchel conjugate f^* of any function f is always convex. Moreover, when f itself is convex and lower semicontinuous⁷, then f^* contains enough information to recover f , in the sense that the biconjugate function f^{**} equals f .

Without going into any details here, one can show the following theorem.

Theorem 3.5 (Characterization of solutions to (3.4)).

A point $x^* \in \mathbb{R}^n$ is a (global) minimizer for (3.4) if and only if there exists $p^* \in \mathbb{R}^m$ such that, for any $\sigma, \tau \neq 0$,

$$x^* \text{ solves } \text{Minimize } \frac{1}{2} \|y - (x^* - \tau A^T p^*)\|_2^2 + \tau f(y) \quad \text{where } y \in \mathbb{R}^n, \quad (3.8a)$$

$$\text{and } p^* \text{ solves } \text{Minimize } \frac{1}{2} \|q - (p^* + \sigma A x^*)\|_2^2 + \sigma g^*(q) \quad \text{where } q \in \mathbb{R}^m. \quad (3.8b)$$

The above theorem provides the idea for an algorithm. As the unknown solutions x^* and p^* appear in the respective objectives for both problems, we can convert (3.8) into a fixed-point algorithm.

Algorithm 3.6 (Primal-dual extragradient algorithm (Chambolle, Pock, 2011)).

1: Set $x^{(k+1)}$ to the solution of the following problem:

$$\text{Minimize } \frac{1}{2} \|y - (x^{(k)} - \tau A^T \bar{p}^{(k)})\|_2^2 + \tau f(y) \quad \text{where } y \in \mathbb{R}^n \quad (3.9)$$

⁷A function $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ is **lower semicontinuous** if $f(x_0) \leq \liminf_{x \rightarrow x_0} f(x)$ holds for all $x_0 \in \mathbb{R}^n$.

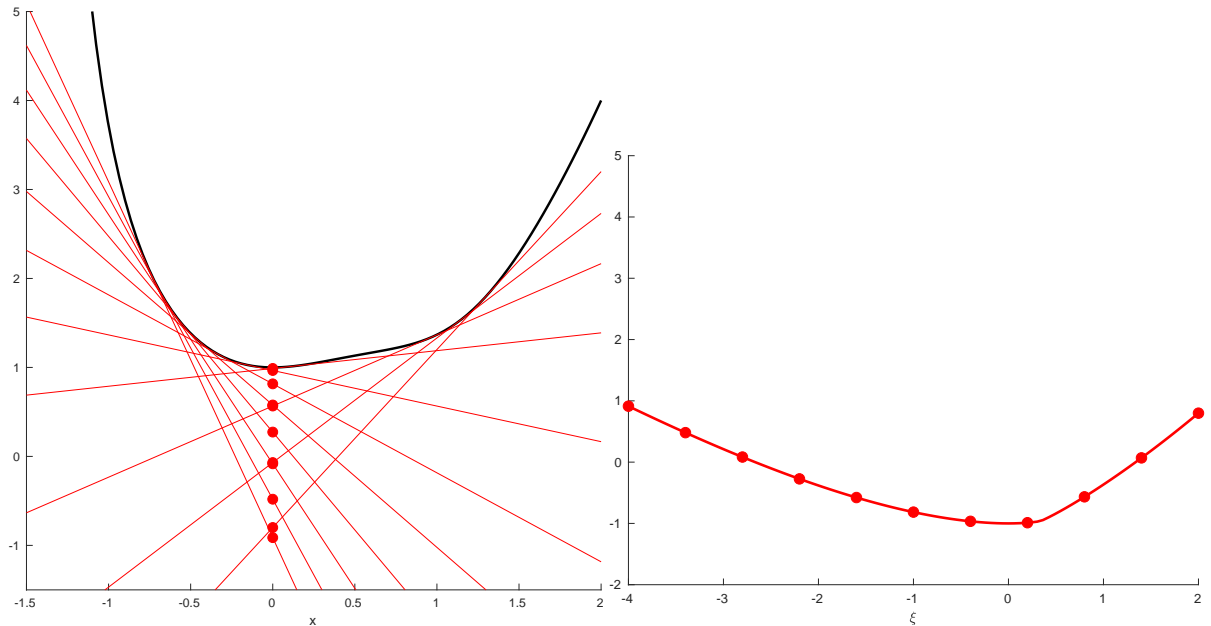


Figure 3.3: Illustration of the construction of the Fenchel conjugate function $f^* : \mathbb{R} \cup \{\pm\infty\}$ (right) of the (convex) function $f(x) = x^2 + \exp(-x^3)$ (left).

2: Set $p^{(k+1)}$ to the solution of the following problem:

$$\text{Minimize } \frac{1}{2} \|q - (p^{(k)} + \sigma Ax^{(k+1)})\|_2^2 + \sigma g^*(q) \quad \text{where } q \in \mathbb{R}^m \quad (3.10)$$

3: Set $\bar{p}^{(k+1)} := 2p^{(k+1)} - p^{(k)}$.

The advantage of such an algorithm is that the solution to the original problem (3.4), involving both f and g , is broken down into the repeated solution to two different problems (3.9) and (3.10). Problem (3.9) requires the minimization of f plus a quadratic distance term, while problem (3.10) requires the minimization of the conjugate g^* plus a quadratic distance term. The solutions to these problems are also known as **proximity (prox) operators**. More precisely, one defines $\text{prox}_{\tau f}(x)$ as the unique solution of the problem “Minimize $\frac{1}{2} \|y - x\|_2^2 + \tau f(y)$.”

For a number of relevant problems, the solutions to problems (3.9) and (3.10) can be explicitly calculated. We consider this now for f and g^* from the image denoising Example 3.3. We begin with $\text{prox}_{\tau f}(x)$, i. e., the unique solution of the problem

$$\text{Minimize } \frac{1}{2} \|y - x\|_2^2 + \tau f(y) = \frac{1}{2} \|y - x\|_2^2 + \frac{\tau}{2} \|y - z\|_2^2 \quad \text{where } y \in \mathbb{R}^{n_1 n_2}.$$

A short calculation shows (**Quiz 3.1**: Can you fill in the details?)

$$\text{prox}_{\tau f}(x) = \frac{x + \tau z}{1 + \tau}, \quad (3.11)$$

i. e., $\text{prox}_{\tau f}(x)$ is simply a weighted average between x and z . On the other hand, we first need to evaluate

$$g^*(\xi) = \sup_{p \in \mathbb{R}^m} \xi^\top p - g(p) = \sup_{p \in \mathbb{R}^m} \xi^\top p - \beta \|p\|_1$$

where $m = n_1(n_2 - 1) + n_2(n_1 - 1)$ is the number of pairs of neighboring pixels (edges in Figure 3.2). We calculate (**Quiz 3.2:** Can you fill in the details?)⁸

$$g^*(\xi) = \begin{cases} 0 & \text{if } \|\xi\|_\infty \leq \beta \\ \infty & \text{if } \|\xi\|_\infty > \beta. \end{cases}$$

Therefore we have that $\text{prox}_{\sigma g^*}(p)$ is the unique solution of the problem

$$\text{Minimize } \frac{1}{2} \|q - p\|_2^2 + \sigma g^*(q) \quad \text{where } q \in \mathbb{R}^m,$$

which is to say that $\text{prox}_{\sigma g^*}(p)$ is the unique solution of

$$\begin{aligned} &\text{Minimize } \frac{1}{2} \|q - p\|_2^2 \quad \text{where } q \in \mathbb{R}^m \\ &\text{subject to } \|q\|_\infty \leq \beta. \end{aligned}$$

A short calculation shows (**Quiz 3.3:** Can you fill in the details?)

$$[\text{prox}_{\sigma g^*}(p)]_i = \frac{\beta p_i}{\max\{\beta, |p_i|\}} = \text{proj}_{[-\beta, \beta]}(p_i), \tag{3.12}$$

which turns out to be independent of σ . This formula means that the values of p are simply clipped elementwise to lie inside $[-\beta, \beta]$.

It can be shown⁹ that Algorithm 3.6 converges to a solution of (3.4) as long as σ, τ are positive numbers satisfying $\sigma \tau \|A\|^2 < 1$, where $\|A\|$ is the largest singular value of A .

We close this section by mentioning what functions a user has to implement in order to run the Chambolle-Pock Algorithm 3.6 for solving (3.4):

proximity operator of τf	$\text{prox}_{\tau f}(x)$
proximity operator of σg^*	$\text{prox}_{\sigma g^*}(p)$
matrix-vector product with A	Ax
matrix-vector product with A^\top	$A^\top p$ redundant if A is available as a matrix

Figure 3.4 shows the result of an image denoising optimization run.

End of Class 2

⁸It can be shown in general that the convex conjugate of a norm is the indicator function of the unit ball w.r.t. the dual norm.

⁹Chambolle, Pock, 2011; Clason, 2017



Figure 3.4: Image denoising model problem (Example 3.3.) Left: original image (dimensions $n_1 = n_2 = 256$). Middle: image with noise. Right: denoising result obtained using 200 iterations of the Chambolle-Pock Algorithm 3.6 with $\sigma = 1$ and $\tau = 1/10$ and weighting parameter $\beta = 0.08$.

§ 4 CONSTRAINED OPTIMIZATION

We recall from § 1 that we are considering constrained optimization problems of the form

$$\left. \begin{array}{ll} \text{Minimize} & f(x) \quad \text{where } x \in \mathbb{R}^n \\ \text{subject to} & g_i(x) \leq 0 \quad \text{with } i = 1, \dots, n_{\text{ineq}} \\ & \text{and } h_j(x) = 0 \quad \text{with } j = 1, \dots, n_{\text{eq}}. \end{array} \right\} \quad (4.1)$$

These problems are known as **nonlinear optimization problems** or **nonlinear programming problems (NLP)**. We also recall the notion of the **feasible set**

$$F := \{x \in \mathbb{R}^n \mid g_i(x) \leq 0 \text{ for all } i = 1, \dots, n_{\text{ineq}} \text{ and } h_j(x) = 0 \text{ for all } j = 1, \dots, n_{\text{eq}}\} \quad (4.2)$$

associated with (4.1).

Example 4.1 (Production example; slightly modified from John E. Beasley's OR notes¹⁰).

A company makes two products (X and Y) using two machines (A and B). Each unit of X that is produced requires 50 minutes processing time on machine A and 30 minutes processing time on machine B. Each unit of Y that is produced requires 24 minutes processing time on machine A and 33 minutes processing time on machine B. At the start of the current week there are 30 units of X and 90 units of Y in stock. Available processing time on machine A is forecast to be 40 hours and on machine B is forecast to be 35 hours. The demand for X in the current week is forecast to be 70 units and for Y is forecast to be 95 units. Company policy is to maximize the combined sum of the units of X and the units of Y in stock at the end of the week.

Formulate this as an optimization problem. What kind of problem is it? What type of solver would you use to solve it?

Note: Spoiler alert: The solution of [Example 4.1](#) is given in [§ 5](#).

Example 4.2 (Formulating non-smooth problems as smooth ones). *Try to formulate the following optimization problems in the form (1.1) with smooth objective and smooth constraints.*

$$\left. \begin{array}{ll} \text{Minimize} & -x_1^2 + x_2 \\ \text{subject to} & \max\{x_1 + x_2, 2x_1 - x_2\} \leq 3 \end{array} \right\} \quad (4.3)$$

$$\left. \begin{array}{ll} \text{Minimize} & -x_1^2 + x_2 \\ \text{subject to} & \min\{x_1 + x_2, 2x_1 - x_2\} \leq 3 \end{array} \right\} \quad (4.4)$$

$$\left. \begin{array}{ll} \text{Minimize} & \max\{-x_1^2 + x_2, x_1^2 + 2x_2\} \\ \text{subject to} & x_1 + x_2 \leq 3 \end{array} \right\} \quad (4.5)$$

$$\left. \begin{array}{ll} \text{Minimize} & | -x_1^2 + x_2 | + | x_1 + x_2^2 | \\ \text{subject to} & x_1 + x_2 \leq 3 \end{array} \right\} \quad (4.6)$$

¹⁰<https://people.brunel.ac.uk/~mastjjb/jeb/or/morelp.html>

We will now look a little bit into the theory for problem (1.1) since it also tells us something about choices we may be facing while modeling the problem.

One can easily see by example (**Quiz 4.1**: Can you give an example?) that the well-known condition $\nabla f(x) = 0$ is no longer a necessary condition for local optimality in case of a constrained problem. It is one of the main conceptual difficulties in constrained optimization to come up with conditions that replace $\nabla f(x) = 0$.

The main idea is as follows. Consider a local minimizer x^* and a sequence $x^{(k)}$ of feasible points approaching it. The elements of the sequence will eventually (i. e., for k large enough) all remain in the neighborhood $U(x^*)$ of local optimality. Consequently,

$$f(x^*) \leq f(x^{(k)})$$

holds for all k sufficiently large. Therefore, we have

$$0 \leq f(x^{(k)}) - f(x^*). \quad (4.7)$$

By Taylor's theorem¹¹, for any $k \in \mathbb{N}$ there exist a number $\xi^{(k)} \in (0, 1)$ such that

$$f(x^{(k)}) = f(x^*) + \nabla f(x^* + \xi^{(k)}(x^{(k)} - x^*))^\top (x^{(k)} - x^*)$$

holds. Plugging this into the inequality (4.7), we obtain

$$0 \leq \nabla f(x^* + \xi^{(k)}(x^{(k)} - x^*))^\top (x^{(k)} - x^*) \quad (4.8)$$

for sufficiently large $k \in \mathbb{N}$. Notice that since $x^{(k)} \rightarrow x^*$ and since $\xi^{(k)}$ is bounded, we must have

$$x^* + \xi^{(k)}(x^{(k)} - x^*) \rightarrow x^*$$

and therefore

$$\nabla f(x^* + \xi^{(k)}(x^{(k)} - x^*)) \rightarrow \nabla f(x^*).$$

However, passing to the limit in (4.8) only yields the meaningless $0 \leq 0$ so we must do something different.

Let us take another sequence $t^{(k)}$ which converges to zero from above: $t^{(k)} \searrow 0$, and let us restrict our attention to sequences $x^{(k)}$ such that the limit

$$d := \lim_{k \rightarrow \infty} \frac{x^{(k)} - x^*}{t^{(k)}}$$

exists. Dividing (4.8) by $t^{(k)}$ gives

$$0 \leq \nabla f(x^* + \xi^{(k)}(x^{(k)} - x^*))^\top \left(\frac{x^{(k)} - x^*}{t^{(k)}} \right)$$

and passing to the limit now yields the informative statement

$$0 \leq \nabla f(x^*)^\top d$$

for all directions $d \in \mathbb{R}^n$ constructed as above. We refer to those as **tangent directions** or **tangent vectors**.

¹¹This requires that the objective f is a continuously differentiable (C^1) function.

Definition 4.3 (Tangent vectors, tangent cone). A vector $d \in \mathbb{R}^n$ is said to be a **tangent direction** or **tangent vector** to the (feasible) set F at the point $x^* \in F$ if there exist sequences $x^{(k)}$ and $t^{(k)}$ such that $x^{(k)} \rightarrow x^*$ and $t^{(k)} \searrow 0$ such that

$$d = \lim_{k \rightarrow \infty} \frac{x^{(k)} - x^*}{t^{(k)}}$$

holds. The set of all tangent vectors to F at x^* is said to be the **tangent cone**¹² to F at x^* and we denote it by $\mathcal{T}_F(x^*)$.

Above we have proved the following

Theorem 4.4 (First-order necessary optimality conditions).

Suppose that x^* is a local minimizer of an optimization problem with feasible set F . Then

$$\nabla f(x^*)^\top d \geq 0 \quad \text{for all } d \in \mathcal{T}_F(x^*). \quad (4.9)$$

The difficulty with this theorem is that the tangent cone $\mathcal{T}_F(x^*)$ is generally hard to characterize and thus impossible to use in an optimization algorithm. Notice also that $\mathcal{T}_F(x^*)$ doesn't even use the description of the feasible set F (4.2) in terms of the inequality constraints g_i or the equality constraints h_j .

We therefore introduce another cone as a candidate to replace $\mathcal{T}_F(x^*)$.

Definition 4.5 (Linearized feasible vector, linearizing cone). A vector $d \in \mathbb{R}^n$ is said to be a **linearized feasible direction** to the (feasible) set F described by (1.2) at the point $x^* \in F$ if it satisfies the following conditions:

$$\nabla g_i(x^*)^\top d \leq 0 \quad \text{for all active indices } i = 1, \dots, m, \quad (4.10a)$$

$$\nabla h_j(x^*)^\top d = 0 \quad \text{for all } j = 1, \dots, p. \quad (4.10b)$$

The set of all linearized feasible directions to F at x^* is said to be the **linearizing cone** to F at x^* and we denote it by $\mathcal{T}_F^{\text{lin}}(x^*)$.

Recall that the active indices at x^* are those satisfying $g_i(x^*) = 0$. Therefore, inequalities that are inactive at x^* do not play a role in the definition.

Remark 4.6 (Tangent and linearizing cones). The tangent cone $\mathcal{T}_F(x^*)$ can be viewed as a geometric approximation to the feasible set near the point x^* . By contrast, the linearizing cone $\mathcal{T}_F^{\text{lin}}(x^*)$ is an attempt to locally describe the feasible set in algebraic terms, by means of equalities and inequalities.

¹²In mathematics, a **cone** in \mathbb{R}^n is any set $K \subseteq \mathbb{R}^n$ such that $x \in K$ implies $\alpha x \in K$ for all $\alpha > 0$. In other words, a point x in the cone entails that the entire ray from the origin through x belongs to the cone as well.

One can show (although we do not do it here) that $\mathcal{T}_F(x^*) \subseteq \mathcal{T}_F^{\text{lin}}(x^*)$ holds at any feasible point x^* . Therefore, we cannot expect that [Theorem 4.4](#) continues to hold if $d \in \mathcal{T}_F(x^*)$ were replaced by $d \in \mathcal{T}_F^{\text{lin}}(x^*)$. In order for that to be possible, we require an extra condition that ensures that $\mathcal{T}_F(x^*) = \mathcal{T}_F^{\text{lin}}(x^*)$. Such a condition guarantees the agreement of the geometric and algebraic descriptions and it is called a **constraint qualification**.

There are various types of constraint qualifications, and we consider only the most simple one.

Definition 4.7 (LICQ). *Suppose that the (feasible) set F is given by (1.2) and $x^* \in F$. We say that x^* satisfies the **linear independence constraint qualification** (in short: **LICQ**) if the following set of vectors is linearly independent:*

$$\{\nabla g_i(x)\}_{i \in \mathcal{A}(x^*)} \cup \{\nabla h_j(x)\}_{j=1}^p.$$

Here $\mathcal{A}(x^*)$ denotes the set of active indices, i. e. $\mathcal{A}(x^*) := \{i = 1, \dots, m \mid g_i(x^*) = 0\}$.

One can show that the LICQ at x^* implies $\mathcal{T}_F(x^*) = \mathcal{T}_F^{\text{lin}}(x^*)$. Provided that x^* is also a local minimizer, we thus get from [Theorem 4.4](#) that

$$\nabla f(x^*)^\top d \geq 0 \quad \text{for all } d \in \mathcal{T}_F^{\text{lin}}(x^*). \quad (4.11)$$

What could be the benefit of this, compared to (4.9)?

By a technique called the **Farkas lemma** we can show that (4.11) can be represented by a set (4.12) of equalities and inequalities, which is amenable to numerical algorithms. We state this result in the following lemma, without proof.

Lemma 4.8 ((4.11) as KKT conditions).

Suppose that the (feasible) set F is given by (4.2) and $x^ \in F$. The following two statements are equivalent.*

(a) (4.11) holds.

(b) There exist vectors $\mu \in \mathbb{R}^m$ and $\lambda \in \mathbb{R}^p$ such that the following set of conditions hold:

$$\nabla f(x^*) + \sum_{i=1}^m \mu_i \nabla g_i(x^*) + \sum_{j=1}^p \lambda_j \nabla h_j(x^*) = 0, \quad (4.12a)$$

$$h(x^*) = 0, \quad (4.12b)$$

$$\mu \geq 0, \quad g(x^*) \leq 0, \quad \mu^\top g(x^*) = 0. \quad (4.12c)$$

The system (4.12) is known as the **Karush-Kuhn-Tucker conditions** (in short: **KKT conditions**) of problem (1.1). The vectors μ and λ are called **Lagrange multipliers** associated with the inequality and equality constraints, respectively. A point x^* which admits associated Lagrange multipliers μ and λ so that (4.12) holds is called a **KKT point**.

Condition (4.12a) is often written in terms of the **Lagrangian** associated with problem (4.1), which is defined as

$$\mathcal{L}(x, \mu, \lambda) := f(x) + \mu^\top g(x) + \lambda^\top h(x). \quad (4.13)$$

We can easily see that (4.12a) is nothing but $\nabla_x \mathcal{L}(x^*, \lambda, \mu) = 0$.

We summarize our findings in the following

Theorem 4.9 (First-order necessary optimality conditions under LICQ, compare Theorem 4.4).

Suppose that x^* is a local minimizer of an optimization problem with feasible set F . Moreover, suppose that the LICQ holds at x^* . Then there exist Lagrange multipliers $\mu \in \mathbb{R}^{n_{\text{ineq}}}$ and $\lambda \in \mathbb{R}^{n_{\text{eq}}}$ such that the KKT conditions (4.12) hold. Moreover, μ and λ are unique.

The uniqueness of the multipliers follows from viewing (4.12a) as a linear system of equations for μ and λ and observing that the LICQ implies that the coefficient matrix has linearly independent columns.

A number of comments concerning (4.12) are in order.

- (a) In the absence of inequality constraints ($n_{\text{ineq}} = 0$), (4.12) is a (generally nonlinear) system of equations only.
- (b) Condition (4.12c) – which comes in through inequality constraints – is called a **complementarity system**. Due to the signs of μ and $g(x^*)$, we can equivalently write $\mu^\top g(x^*) = 0$ componentwise as $\mu_i g_i(x^*) = 0$ for all $i = 1, n_{\text{ineq}}$. Consequently, Lagrange multipliers pertaining to inactive inequality constraints must be zero. On the other hand, a strictly positive multiplier $\mu_i > 0$ indicates that the associated constraint must be active, i. e., $g_i(x^*) = 0$.

We can ask the question how “likely” it is for the LICQ to hold for an optimization problem at hand. First of all, the very Definition 4.7 implies that for the LICQ to hold there cannot be too many active inequality and equality constraints in the problem. As soon as the combined number of active inequality and equality constraints at a point x^* exceeds n (the dimension of the optimization variable), the LICQ cannot hold at x^* .

What can we learn from this? When we are modeling an optimization problem, the observation above tells us not to unconsciously throw in additional constraints (which may not be needed) at will. The following example gives us additional hints about things to avoid in modeling constraints:

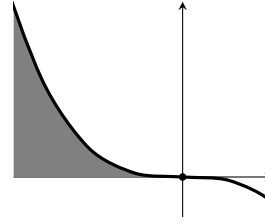
Example 4.10 (LICQ holding or not holding?). Consider the following formulations of a (trivial) optimization problem with variable $x \in \mathbb{R}$, and decide whether or not the LICQ holds at the (obvious) solution:

$$\left\{ \begin{array}{l} \text{Minimize } x^2 \\ \text{subject to } x = 0 \end{array} \right\} \quad \left\{ \begin{array}{l} \text{Minimize } x^2 \\ \text{subject to } x^2 = 0 \end{array} \right\} \quad \left\{ \begin{array}{l} \text{Minimize } x^2 \\ \text{subject to } x^2 \leq 0 \end{array} \right\} \quad \left\{ \begin{array}{l} \text{Minimize } x^2 \\ \text{subject to } x \leq 0 \\ \text{and } x \geq 0 \end{array} \right\}$$

The following example shows that the failure of constraint qualifications to hold can indeed be problematic in practice:

Example 4.11 (Minimizer where the KKT conditions fail to hold). Consider the problem

$$\begin{aligned} \text{Minimize} \quad & -x_1 \quad \text{where } x \in \mathbb{R}^2 \\ \text{subject to} \quad & x_2 + x_1^3 \leq 0 \\ \text{and} \quad & -x_2 \leq 0 \end{aligned}$$



Obviously, $x^* = (0, 0)^\top$ is the unique global minimizer, and there are no further local minimizers. Since both constraints are active at x^* , the KKT conditions (4.12) boil down to

$$\begin{pmatrix} -1 \\ 0 \end{pmatrix} + \mu_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \mu_2 \begin{pmatrix} 0 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Clearly, it is impossible to satisfy them with any $\mu_1, \mu_2 \geq 0$. Hence in this example, the minimizer x^* fails to be a KKT point.

The observation in the previous example is noteworthy since essentially all numerical algorithms for the solution of smooth, constrained problems (4.1) are indeed looking for KKT points.

Example 4.12 (Rosenbrock's post office problem). A simple test example for nonlinear programming is Rosenbrock's post office problem:

$$\begin{aligned} \text{Minimize} \quad & -x_1 x_2 x_3 \quad \text{where } x \in \mathbb{R}^3 \\ \text{subject to} \quad & x_1 + 2x_2 + 2x_3 - 72 \leq 0 \\ \text{and} \quad & x_1 + 2x_2 + 2x_3 \geq 0 \\ \text{and} \quad & 0 \leq x_i \leq 42, \quad i = 1, 2, 3 \end{aligned}$$

This problem has only linear constraints but a nonlinear objective. The global minimizer is $x^* = (24, 12, 12)^\top$.

As we discussed for unconstrained problems in § 2, most optimization algorithms proceed in a sequential way and break down problem (4.1) into a sequence of simpler problems. We will consider here only the class of **sequential quadratic programming (SQP) methods**, where these simpler problems are **quadratic programming (QP) problems**. A QP is a problem of type (4.2) whose objective is a quadratic polynomial and all of whose constraints are linear. In an SQP method, the QP formed at iterate $x^{(k)}$ is of the form

$$\left. \begin{aligned} \text{Minimize} \quad & q^{(k)}(x) := \underbrace{\mathcal{L}(x^{(k)}, \mu^{(k)}, \lambda^{(k)})}_{\text{constant term}} + \underbrace{\mathcal{L}_x(x^{(k)}, \mu^{(k)}, \lambda^{(k)})(x - x^{(k)})}_{\text{linear term}} \\ & + \underbrace{\frac{1}{2}(x - x^{(k)})^\top H^{(k)}(x - x^{(k)})}_{\text{quadratic term}} \quad \text{where } x \in \mathbb{R}^n \\ \text{subject to} \quad & g_i(x^{(k)}) + g'_i(x^{(k)})(x - x^{(k)}) \leq 0 \quad \text{with } i = 1, \dots, n_{\text{ineq}} \\ \text{and} \quad & h_j(x^{(k)}) + h'_j(x^{(k)})(x - x^{(k)}) = 0 \quad \text{with } j = 1, \dots, n_{\text{eq}}. \end{aligned} \right\} \quad (4.14)$$

In comparison to (2.2), the objective is now a quadratic model of the Lagrangian, rather than a quadratic model of the objective. Depending on how the Hessians $H^{(k)}$ of the model are formed, one can distinguish between quasi-Newton and Newton-type SQP methods. As was the case in § 2, precautions need to be taken in order to achieve good convergence properties regardless of the quality of the initial guess provided by the user. As before, there are line-search as well as trust-region but also filter SQP methods around for this purpose.

Many other details need to be taken into account in order to obtain a robust SQP solver, including potentially infeasible QP subproblems, update of the Lagrange multipliers, stopping criteria, deterioration of superlinear convergence due to nonlinear constraints, failure of constraint qualifications etc.

We close this section by mentioning what functions a user has to implement in order to run an SQP for solving a constrained problem (4.1):

objective	$f(x)$	
derivative	$f'(x)$	optional (finite difference fallback)
2nd derivative	$f''(x)$ or $f''(x) d$	optional, only for Newton SQP methods
constraints	$g(x), h(x)$	
Jacobian	$g'(x), h'(x)$	optional (finite difference fallback)
2nd derivative	$g''_i(x), h''_j(x)$ or $g''_i(x) d, h''_j(x) d$	optional, only for Newton SQP methods

End of Class 3

§ 5 SOLUTIONS

§ 5.1 SOLUTION OF EXAMPLE 2.2 (ROSENBROCK) USING CASADI (PROBLEM-BASED)

CASADI provides algorithmic differentiation for user-defined functions. We demonstrate below the description of the problem in a suitable form for CASADI's `nlpsol` solver interface.

```
# This code solves an unconstrained optimization problem with the Rosenbrock
# function using CasADI's python bindings. For lack of a solver specializing
# in unconstrained optimization solver, the solver of choice is nlpsol with
# ipopt, see https://web.casadi.org/docs/#nonlinear-programming.

# Resolve the dependencies.
from casadi import *
import matplotlib.pyplot as plt
import numpy as np

# Create a 1-by-1 optimization variable (termed 'x') and another 1-by-1
# optimization variable (termed 'y').
# SX is the class in CasADI representing matrices composed of symbolic expressions.
# There is also MX for more complex expressions.
x = SX.sym('x', 1, 1)
y = SX.sym('y', 1, 1)

# Set the fixed problem parameters.
a = DM(1)
b = DM(100)

# Setup the objective.
f = (a - x)**2 + b * (y - x**2)**2

# Setup the solver.
problem = {'x': vertcat(x,y), 'f': f}
solver = nlpsol('production', 'ipopt', problem)

# Call the solver.
result = solver()

# Show the solution.
print()
print(result['x'])
```

§ 5.2 SOLUTION OF EXAMPLE 2.2 (ROSENBROCK) USING CASADI (OPTI INTERFACE-BASED)

CasADi offers the Opti interface, which allows a simplified description and solution of optimization problems.

```
# This code solves an unconstrained optimization problem with the Rosenbrock
# function using CasADi's python bindings, and CasADi's Opti interface for
# modeling of (un)constrained optimization problems.

# Resolve the dependencies.
from casadi import *

# Create an instance of the Opti interface.
opti = Opti()

# Create a 1-by-1 optimization variable (termed 'x') and another 1-by-1
# optimization variable (termed 'y').
x = opti.variable()
y = opti.variable()

# Set the fixed problem parameters.
a = opti.parameter()
b = opti.parameter()
opti.set_value(a, 1)
opti.set_value(b, 100)

# Setup the objective.
f = (a - x)**2 + b * (y - x**2)**2
opti.minimize(f)

# Setup the solver.
opti.solver('ipopt')

# Call the solver.
result = opti.solve()

# Show the solution.
print()
print(result.value(x), result.value(y))
```

§ 5.3 SOLUTION OF EXAMPLE 2.3 (AIRPLANE) USING CASADI (PROBLEM-BASED)

Apparently CasADi has no interface for solving least-squares problems. Therefore, we resort to treating the problem as a generic unconstrained optimization problem. Consequently — in contrast to using a dedicated least-squares solver — we have to implement the objective $f(x) = \frac{1}{2} \sum_{i=1}^M [r_i(x)]^2$ instead of the residual function $r(x)$.

```
# This code solves a parameter estimation problem for a given model, based on a
# set of measurement pairs, using CasADi's python bindings for modeling of
# (un)constrained optimization problems.

# Resolve the dependencies.
from casadi import *
import matplotlib.pyplot as plt

# Create the unknown parameters as optimization variables.
xy = MX.sym('xy', 2, 1)

# Specify the pairs of measurements.
measurements = [
    ([ 8,  6], 38),
    ([-3, -3], 220),
    ([ 1,  0], 222),
    ([ 8, -3], 300),
]

# Declare the beacon position as a variable.
beacon = MX.sym('beacon position', 2, 1)

# Specify the model function.
angle = fmod(atan2(beacon[1] - xy[1], beacon[0] - xy[0]) + 2*pi, 2*pi) * 180 / pi
model = Function('airplane', [xy, beacon], [angle], ['xy', 'beacon'], ['angle'])

# Assemble the least-squares objective function.
f = 0
for i in range(len(measurements)):
    residual = model(xy = xy, beacon = measurements[i][0])['angle'] - measurements[i][1]
    f = f + residual**2
f = 0.5 * f

# Setup the initial guess.
x0 = [0, 0]

# Setup the solver.
problem = {'x': xy, 'f': f}
solver = nlpsol('airplane', 'ipopt', problem)

# Call the solver.
result = solver(x0 = x0)

# Show the solution.
print()
print(result['x'])
```


§ 5.4 SOLUTION OF EXAMPLE 4.1 (PRODUCTION) USING MATLAB'S OPTIMIZATION TOOLBOX (SOLVER-BASED)

The classical use of MATLAB's optimization toolbox required the user to model their optimization problem in a format suitable for the respective solver to be used. Since [Example 4.1](#) is a linear optimization problem, `linprog` is the solver of choice.

```
% This code solves the production example problem using Matlab's optimization
% toolbox. The problem is a linear optimization problem. The solver of choice is
% linprog from Matlab's optimization toolbox, see
% https://de.mathworks.com/help/optim/ug/linprog.html.

% Setup the cost vector describing the objective c'*x.
c = [-1; -1];

% Setup the matrix A and right hand side b for the inequality constraint
% A * x <= b.
A = [50, 24; 30, 33];
b = [2400; 2100];

% Setup the matrix Aeq and right hand side beq for the (void) equality constraint
% Aeq * x = beq.
Aeq = [];
beq = [];

% Setup the lower and upper bounds.
lb = [36; 5];
ub = [];

% Call the solver.
[x, fval] = linprog(c, A, b, Aeq, beq, lb, ub);

% Show the solution.
disp(x)
```

§ 5.5 SOLUTION OF EXAMPLE 4.1 (PRODUCTION) USING MATLAB'S OPTIMIZATION TOOLBOX (PROBLEM-BASED)

With more recent versions of MATLAB's optimization toolbox, the user may formulate their problem in a more natural way, and independently of the solver. The solver will be picked automatically depending on the problem characteristics.

```
% This code solves the production example problem using Matlab's optimization
% toolbox. The problem is modeled using the problem-based approach, which
% allows for a more natural problem formulation and automatic selection of a
% suitable solver, see
% https://de.mathworks.com/help/optim/problem-based-basics.html.

% Create a maximization-type problem.
problem = optimproblem('ObjectiveSense', 'maximize', 'Description', 'production');

% Create a 2-by-1 optimization variable (termed 'x').
x = optimvar('x', 2, 1, 'LowerBound', [36; 5]);

% Setup the objective.
problem.Objective = x(1) + x(2);

% Setup the (named) inequality constraints.
problem.Constraints.timeConstraintA = 50 * x(1) + 24 * x(2) <= 2400;
problem.Constraints.timeConstraintB = 30 * x(1) + 33 * x(2) <= 2100;

% Review the problem.
show(problem)

% Call the solver.
result = solve(problem);

% Show the solution.
disp(result.x)
```

§ 5.6 SOLUTION OF EXAMPLE 4.1 (PRODUCTION) USING `scipy.optimize`

Here we demonstrate the solution of [Example 4.1](#) with `scipy.optimize.linprog`. We need to present the problem in a suitable form to the solver.

```
# This code solves the production example problem using the scipy.optimize
# package. The problem is a linear optimization problem. The solver of choice is
# scipy.optimize.linprog, see
# https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.linprog.html.

# Resolve the dependencies.
from scipy.optimize import linprog

# Setup the cost vector describing the objective  $c @ x$ .
c = [-1, -1]

# Setup the matrix A and right hand side b for the inequality constraint
#  $A @ x \leq b$ .
A = [[50, 24], [30, 33]]
b = [2400, 2100]

# Setup the lower and upper bounds.
bounds = [(36, None), (5, None)]

# Call the solver.
result = linprog(c, A_ub = A, b_ub = b, bounds = bounds)

# Show the solution.
print(result['x'])
```

§ 5.7 SOLUTION OF EXAMPLE 4.1 (PRODUCTION) USING CASADI (PROBLEM-BASED)

CASADI provides algorithmic differentiation for user-defined functions (although this is not needed for linear optimization problems). We demonstrate below the description of the problem in a suitable form for CASADI's `nlpso1` solver interface.

```
# This code solves the production example problem using CasADI's python
# bindings. The problem is a linear optimization problem. For lack of a linear
# optimization solver, the solver of choice is nlpso1 with ipopt, see
# https://web.casadi.org/docs/#nonlinear-programming.

# Resolve the dependencies.
from casadi import *

# Create a 2-by-1 optimization variable (termed 'x').
# SX is the class in CasADI representing matrices composed of symbolic expressions.
# There is also MX for more complex expressions.
x = SX.sym('x', 2, 1)

# Setup the objective.
f = - x[0] - x[1]

# Setup the constraint function g representing the inequality constraint
# A @ x <= b.
# DM is the class in CasADI representing matrices with numerical values.
A = DM([[50, 24], [30, 33]])
g = A @ x

# Setup the upper bound for the constraint function g.
b = DM([2400, 2100])

# Setup the lower bound for the variable x.
l = DM([36, 5])

# Setup the solver.
problem = {'x': x, 'f': f, 'g': g}
solver = nlpso1('production', 'ipopt', problem)

# Call the solver, providing the upper bounds for the inequality constraint
# function g and the lower bound for the optimization variable.
result = solver(ubg = b, lbx = l)

# Show the solution.
print()
print(result['x'])
```

§ 5.8 SOLUTION OF EXAMPLE 4.1 (PRODUCTION) USING CASADI (OPTI INTERFACE-BASED)

CasADi offers the Opti interface, which allows a simplified description and solution of optimization problems.

```
# This code solves the production example problem using CasADi's python
# bindings, and CasADi's Opti interface for modeling of (un)constrained
# optimization problems.

# Resolve the dependencies.
from casadi import *

# Create an instance of the Opti interface.
opti = Opti()

# Create a 2-by-1 optimization variable (termed 'x').
x = opti.variable(2)

# Setup the objective.
opti.minimize(- x[0] - x[1])

# Setup the constraint function representing the inequality constraint
# A @ x <= b.
# DM is the class in CasADi representing matrices with numerical values.
A = DM([[50, 24], [30, 33]])
b = DM([2400, 2100])
opti.subject_to(A@x <= b)

# Setup the lower bound for the variable x.
l = DM([36, 5])
opti.subject_to(x >= l)

# Setup the solver.
opti.solver('ipopt')

# Call the solver.
result = opti.solve()

# Show the solution.
print()
print(result.value(x))
```

§ 5.9 SOLUTION OF EXAMPLE 4.12 (POST-OFFICE) USING CASADI (OPTI INTERFACE-BASED)

CASADI offers the `Opti` interface, which allows a simplified description and solution of optimization problems.

```
# This code solves Rosenbrock's post-office example problem using CasADI's
# python bindings, and CasADI's Opti interface for modeling of (un)constrained
# optimization problems.

# Resolve the dependencies.
from casadi import *

# Create an instance of the Opti interface.
opti = Opti()

# Create a 3-by-1 optimization variable (termed 'x').
x = opti.variable(3)

# Setup the objective.
opti.minimize(- x[0] * x[1] * x[2])

# Setup the constraint functions.
opti.subject_to(x[0] + 2 * x[1] + 2 * x[2] - 72 <= 0)
opti.subject_to(x[0] + 2 * x[1] + 2 * x[2] >= 0)
opti.subject_to(opti.bounded(0, x, 42))

# Setup the solver.
opti.solver('ipopt')

# Call the solver.
result = opti.solve()

# Show the solution.
print()
print(result.value(x))
```

Bibliography

- Andersson, J. A. E.; J. Gillis; G. Horn; J. B. Rawlings; M. Diehl (2018). “CasADi: a software framework for nonlinear optimization and optimal control”. *Mathematical Programming Computation* 11.1, pp. 1–36. DOI: [10.1007/s12532-018-0139-4](https://doi.org/10.1007/s12532-018-0139-4).
- Chambolle, A.; T. Pock (2011). “A first-order primal-dual algorithm for convex problems with applications to imaging”. *Journal of Mathematical Imaging and Vision* 40.1, pp. 120–145. DOI: [10.1007/s10851-010-0251-1](https://doi.org/10.1007/s10851-010-0251-1).
- Clason, C. (2017). *Nonsmooth analysis and optimization*. arXiv: [1708.04180](https://arxiv.org/abs/1708.04180).
- Rudin, L. I.; S. Osher; E. Fatemi (1992). “Nonlinear total variation based noise removal algorithms”. *Physica D* 60.1–4, pp. 259–268. DOI: [10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F).