Norwegian University of Science and Technology
Department of Mathematical Sciences

TMA4180
Optimisation I
Spring 2019

Exercise set "Project"

1 Segmentation

In this project, we will look at an optimisation problem in data segmentation and model fitting. We assume that we are given a collection of observations or data points $z_i \in \mathbb{R}^d$, $i = 1, \ldots, m$, together with labels $w_i \in \{a, b\}$. Our goal is to separate the data points with the different labels using second order polynomials, in particular ellipsoids and circles.

In the first case we are considering, we assume that the data points with label a lie mainly in the interior of an ellipsoid in $\mathbb{R}^d$, whereas the data points lie mainly in the exterior of that same ellipsoid. However, because of noise (or modelling) errors, the boundary of the ellipsoid does not separate the two point clouds strictly (see Figure 1). Later on, we will (possibly) look into the situation, where we want to find two distinct ellipsoids separating the two sets of data points.

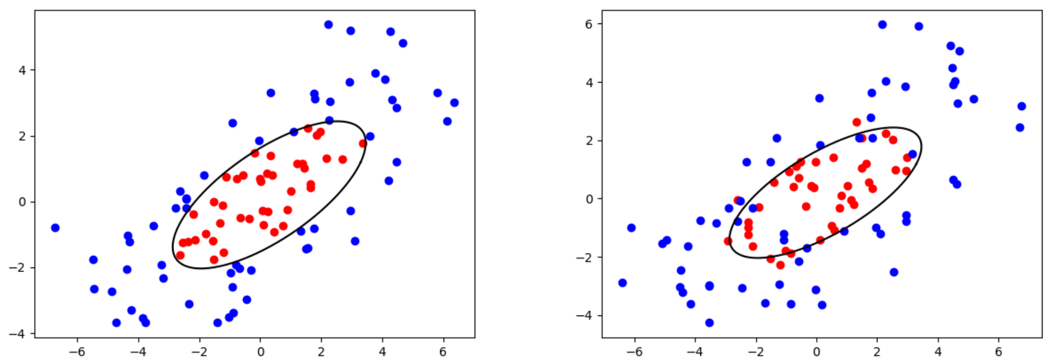

Figure 1: *Left:* Ideal situation, where data points are well separated by an ellipse. *Right:* Typical noisy case, where an exact separation is no longer possible.

2 Second order hypersurfaces

In order to formulate an optimisation problem for the computation of the separating ellipsoids, we have to discuss briefly how ellipsoids (and similar hypersurfaces) can be described.

- A general ellipsoid in $\mathbb{R}^d$ can be defined by means of a positive definite and symmetric matrix $A \in \mathbb{R}^{d\times d}$ and a point $c \in \mathbb{R}^d$ as
$$S_{A,c} = \{x \in \mathbb{R}^d : (x-c)^T A(x-c) \le 1\}.$$
The point c is the center of the ellipsoid, the directions of the semi-axes are the eigendirections of A, and their lengths are the square roots of the inverses of the corresponding eigenvalues of A.

- An alternative description of an ellipsoid uses again a positive definite and symmetric matrix $A \in \mathbb{R}^{d\times d}$, a vector $b \in \mathbb{R}^d$ and a number $\gamma \in \mathbb{R}$. Here we can define
$$E_{A,b,\gamma} = \{x \in \mathbb{R}^d : x^T Ax - x^T b - \gamma \le 0\}.$$
This corresponds to the description as $S_{A,c}$ with $b = 2Ac$ and $\gamma = 1 - c^T Ac$. Note, however, that this description is not unique (if we multiply A, b, and γ with the same positive scalar, the ellipse does not change), and $E_{A,b,\gamma}$ is empty for certain combinations of A, b, and γ.

- If we drop the assumption that A is positive definite (but still assume that is is symmetric), we can use these formulas to describe more general second order shapes like hyperpoloids or paraboloids. In the case of $E_{A,b,\gamma}$, we also obtain halfspaces (with $A = 0$).

- Although we cannot generate all possible ellipsoids (or more general second order surfaces), it can make sense to restrict ourselves to the sets
$$E_{A,b} := E_{A,b,1} = \{x \in \mathbb{R}^d : x^T Ax - x^T b \le 1\}.$$
In this case, each set $E_{A,b}$ is non-empty and in particular contains the point $x = 0$. Moreover, two sets $E_{A,b}$ and $E_{C,d}$ are equal if and only if $A = C$ and $b = d$. However, this only makes sense if we know that we can restrict ourselves to situations where $0 \in E_{A,b}$.

3 Symmetric and positive semi-definite matrices

We denote in the following by
$$\mathrm{Sym}_d := \{A \in \mathbb{R}^{d\times d} : A^T = A\}$$
the set of all symmetric $(d \times d)$-matrices. Note that we can identify a symmetric matrix $A \in \mathrm{Sym}_d$ with a $d(d+1)/2$-dimensional vector containing the entries of A on and above the diagonal. In particular, for $d = 2$ we can identify
$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{12} & A_{22} \end{pmatrix} \in \mathrm{Sym}_2$$
with the vector
$$a = \begin{pmatrix} A_{11} \\ A_{12} \\ A_{22} \end{pmatrix} \in \mathbb{R}^3.$$
Moreover, such a matrix A is positive semi-definite if and only if
$$A_{11} \ge 0, \qquad A_{22} \ge 0, \qquad \text{and} \qquad \det A = A_{11}A_{22} - A_{12}^2 \ge 0,$$
and positive definite, if and only if all these inequalities are strict.

4 Unconstrained optimisation

We start with the simplest situation depicted in Figure 1, where the different points are separated by a single hypersurface. We will consider two different models, the first based on the sets $S_{A,c}$, the second on the sets $E_{A,b}$. For simplicity, we will drop the assumption of positive definiteness of A for the moment. The goal is to find a set R of the form $R = S_{A,c}$ or $R = E_{A,b}$, such that the conditions $z_i \in R$ for $w_i = a$ and $z_i \notin R$ for $w_i = b$ are satisfied as good as possible. That is, the points with label a should be contained in R, whereas the points with label b should be outside of R.

4.1 First model

Given a data point $z_i \in \mathbb{R}^d$, we define the mapping $r_i \colon \operatorname{Sym}_d \times \mathbb{R}^d \to \mathbb{R}$,

$$r_i(A, c) := (z_i - c)^T A(z_i - c) - 1.$$

That is, a point z_i is contained in $S_{A,c}$ if $r_i(A, c) \leq 0$, and is not contained in $S_{A,c}$ if $r_i(A, c) > 0$. In order to find the optimal set $S_{A,c}$ it makes therefore sense to solve the problem

$$\min_{\substack{A \in \operatorname{Sym}_d \\ c \in \mathbb{R}^d}} f_1(A, c), \tag{1}$$

where

$$f_1(A, c) = \sum_{w_i = a} \max\{r_i(A, c), 0\}^2 + \sum_{w_i = b} \min\{r_i(A, c), 0\}^2.$$

That is, points z_i with label $w_i = a$ do not contribute to the function value if $z_i \in S_{A,c}$ but contribute with the squared residuum $r_i(A, c)$ if $z_i \notin S_{A,c}$, and points z_i with label $w_i = c$ contribute quadratically if $z_i \in S_{A,c}$. In other words, each misclassified point provides a contribution to the function value equal to the squared residual.

4.2 Second model

We use the same approach as in Section 4.1, but based on the sets $E_{A,b}$ instead of $S_{A,c}$: Given a data point $z_i \in \mathbb{R}^d$, we define the mapping $\tilde{r}_i \colon \operatorname{Sym}_d \times \mathbb{R}^d \to \mathbb{R}$,

$$\tilde{r}_i(A, b) := z_i^T A z_i - z_i^T b - 1.$$

In order to segment the data points $z_i \in \mathbb{R}^d$, we then consider the optimisation problem

$$\min_{\substack{A \in \operatorname{Sym}_d \\ b \in \mathbb{R}^d}} f_2(A, b), \tag{2}$$

where

$$f_2(A, b) = \sum_{w_i = a} \max\{\tilde{r}_i(A, b), 0\}^2 + \sum_{w_i = b} \min\{\tilde{r}_i(A, b), 0\}^2.$$

Exemplary results of the problems of minimising f_1 and f_2 are shown in Figure 2.

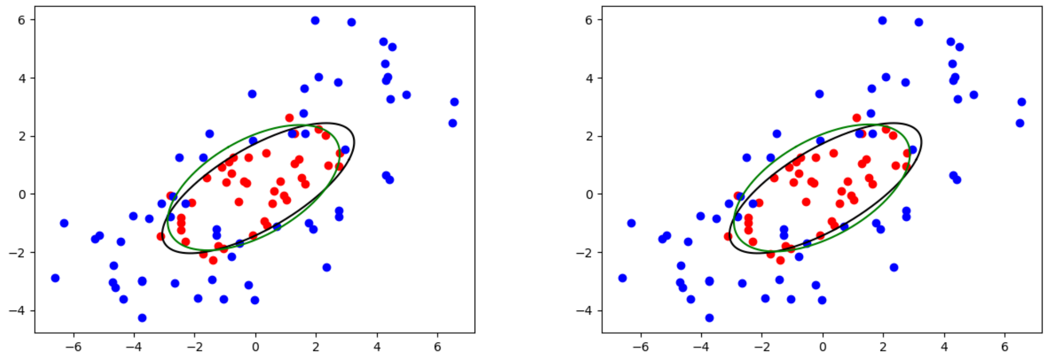

Figure 2: *Left:* Minimiser of f_1. *Right:* Minimiser of f_2. The green ellipse depicts the solution (that is, the boundary of the ellipse $S_{A,c}$ or $E_{A,b}$), whereas the "true ellipse", which was used to generate the data, is depicted in black. Although the results are very close, one can see small differences that result from the different definitions of the residua r_i and $\tilde{r}_i$.

4.3 Problems

1. Show that the mappings f_1 and f_2 defined above are C^1, but typically not C^2.

2. Show that the mapping f_2 is convex and that the gradient ∇f_2 is Lipschitz continuous.[1]

3. Discuss the existence and uniqueness of solutions of (1) and (2).

4. Implement numerical solution methods in dimension $d = 2$ for the problems (1) and (2).

5. Construct reasonable test cases and compare the behaviour of the optimisation algorithms you implemented for the two different functionals f_1 and f_2 on these test cases.

4.4 Test cases

One method for constructing test cases is to start by choosing parameters (A, c) or (A, b). Then one can generate a random sample of data points z_i and assign to them labels

[1] It can help to show first that a composition of a convex function and an affine function is again convex.

according to whether they lie within the set $S_{A,c}$ (or $E_{A,b}$) or not. Afterwards it is possible to perturb the data points z_i slightly by adding a small random vector to each of the points (without changing the label). An alternative method for modelling classification errors is to change a small number of labels. In addition, one can try to approximate other sets, such as rectangles or other polytopes, using model sets $S_{A,c}$ or $E_{A,b}$ by minimising f_1 or f_2. Similarly to the previous case, one can generate some random points, but then assign the labels based on whether the data points are contained in the rectangle/polytope.

Experiment with different problem types/sizes (in terms of the number of data points) in order to see how the algorithms you have implemented perform in practice. The data can, for instance, be visualized as a scatter plot, coloured according to the labels w_i. The computed model sets can then be vizualized as a zero-level contour of the function

$$z \mapsto (z-c)^T A(z-c) - 1$$

or

$$z \mapsto z^T Az - z^T b - 1.$$

5 Positivity constraints

5.1 Setting

The methods that we have formulated until now do not take into account that the reconstructed hypersurfaces should describe ellipsoids, as we have ignored the requirement that the matrix A should be positive definite. As a consequence, the hypersurface that is defined by the parameters (A, c) or (A, b) need not be an ellipsiod, but can also be a hyperboloid, paraboloid, or even a strip or half-space. Depending on the context, this can be either seen as a benefit or a problem. In the following, we will take the latter point of view and thus try to restrict our minimisation problems to positive definite matrices A.

In general, optimisation with positive-definiteness constraints can be quite challenging, as the smallest (and largest) eigenvalue of a matrix depends in a non-smooth way on its entries. Thus we will restrict ourselves to the two-dimensional case, where bounds on the eigenvalues are easier to implement. Moreover, we will try to implement both upper and lower bounds on the eigenvalues of A in order to exclude ellipses with eccentricities close to zero.

Assume now that $A \in \mathbb{R}^{2\times 2}$ is a symmetric matrix with eigenvalues λ_1, $\lambda_2 \in \mathbb{R}$. Then A is positive definite if and only if $\lambda_i > 0$, $i = 1, 2$. Alternatively, we have that A is positive definite if $A_{11} > 0$, $A_{22} > 0$, and $\det A = A_{11}A_{22} - A_{12}^2 > 0$. One possible approach is therefore, to solve one of the optimisation problems

$$\min_{\substack{A\in \mathrm{Sym}_2 \\ c\in\mathbb{R}^2}} f_1(A, c) \qquad \text{or} \qquad \min_{\substack{A\in \mathrm{Sym}_2 \\ b\in\mathbb{R}^2}} f_2(A, b) \tag{3}$$

subject to the constraints

$$\left.\begin{aligned} \gamma_1 \le A_{11} \le \gamma_2, \\ \gamma_1 \le A_{22} \le \gamma_2, \\ (A_{11}A_{22})^{1/2} \ge (\gamma_1^2 + A_{12}^2)^{1/2} \end{aligned}\right\} \tag{4}$$

for some fixed parameters

$$0 < \gamma_1 < \gamma_2 < \infty.$$

5.2 Problems

All that follows builds upon the problems in Section 4.3 dealing with the unconstrained case. Make sure that you have a good grasp of the theory and numerics of that unconstrained situation (and also a working numerical algorithm) before you continue with the following problems.

6 Show that the function

$$(A_{11}, A_{12}, A_{22}) \mapsto (A_{11}A_{22})^{1/2} - (\gamma_1^2 + A_{12}^2)^{1/2}$$

is C^2 and concave on the set

$$\{(A_{11}, A_{12}, A_{22}) \in \mathbb{R}^3 : A_{11} > 0 \text{ and } A_{22} > 0\}.$$

Conclude that the constraints (4) describe a convex and closed set.

7 Show that the constraints (4) satisfy Slater's constraint qualification.[2]

8 Formulate the KKT conditions for the problems in (3) with constraints (4). Discuss whether these conditions are necessary or sufficient optimality conditions.

9 Implement a numerical optimisation method for solving the problems in (3) with constraints (4).

10 Discuss the behaviour of your method on different test cases. Include also cases in your study, where the "true solution" (that is, the model from which the data points have been generated) does not satisfy the constraints.

6 Separating several sets

6.1 Setting

We now expand/modify our problem to a setting where we are given data points z_i with three different possible labels $w_i \in \{a, b, c\}$, and we assume that the points with label a

[2]Recall the definition of Slater's constraint qualification:
If the set Ω is described by the inequality constraints $c_i(x) \geq 0$, $i \in \mathcal{I}$, then Slater's constraint qualification is satisfied if all the functions c_i are concave and there exists a point $\hat{x}$ with $c_i(\hat{x}) > 0$, $i \in \mathcal{I}$. In such a situation we have at every point $x \in \Omega$ that $T_\Omega(x) = \mathcal{F}(x)$.

belong to some set R_1, the points with label b to some set R_2, and the points with label c to none of these sets. More specifically, we consider the situation where the sets R_i are (disjoint) balls

$$R_i = \{\|z - c_i\|^2 \le \rho_i^2\}$$

(see Figure 3).

In a similar manner as above, we can then define the residuum

$$r_i(c, \rho) := \|z_i - c\|^2 - \rho^2$$

and the function

$$f_3(c, \rho; d, \sigma) := \sum_{w_i = a} \max\{r_i(c, \rho), 0\}^2 + \sum_{w_i \in \{b,c\}} \min\{r_i(c, \rho), 0\}^2 + \sum_{w_i = b} \max\{r_i(d, \sigma), 0\}^2 + \sum_{w_i \in \{a,c\}} \min\{r_i(d, \sigma), 0\}^2.$$

Then a possible segmentation of the data points can be obtained by solving the constrained optimisation problem

$$\min_{(c,\rho;d,\sigma) \in \mathbb{R}^6} f_3(c, \rho; d, \sigma)$$

subject to the constraints

$$\rho \ge 0, \qquad \sigma \ge 0, \qquad \|c - d\| \ge \rho + \sigma. \tag{5}$$

The last inequality here simply means that the two balls we are constructing are disjoint.

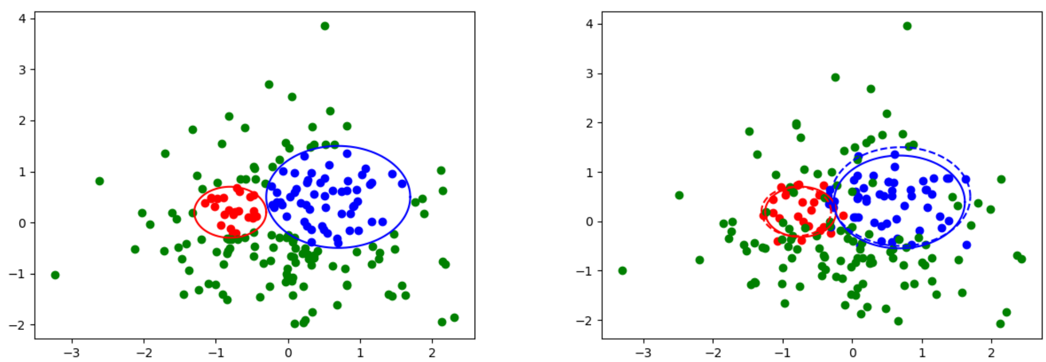

Figure 3: *Left:* Noise-free data set. *Right:* Noisy data set and possible reconstruction. Dashed lines: "True solutions." Solid lines: Numerical solution of the optimisation problem.

6.2 Problems

Only start working on the following problems after you have dealt with the problems in Section 5.2. It is better to submit a good project report where only problems 1–10 are treated, than a report with incomplete answers to all problems.

11 Discuss the convexity of the function f_3 and the set given by the constraints in (5).

12 Formulate the KKT conditions for the problem of minimising f_3 subject to the constraints in (5). Discuss whether these conditions are necessary or sufficient optimality conditions.

13 Implement a numerical optimisation method for solving the problem of minimising f_3 subject to the constraints in (5), and discuss the behaviour of that method for different test cases.

The last problem is intended as an open ended and potentially challenging problem for those students who have too much spare time on their hands. It is completely optional and not required, feel free to have fun with it, though (you do not need to discuss this problems for obtaining full marks in the project).

14 Investigate how this model can be generalised to the situation where the sets R_i one wants to reconstruct are disjoint ellipses $S_{A,c}$ and $S_{B,d}$, and propose a numerical solution for the resulting optimisation problem.