MATH-329 Nonlinear optimization Homework 2: Constrained optimization

Anya-Aurore Mauron & Benoît Müller

October 4, 2023

Part 1 : Projections to cones and stopping criteria in constrained optimization.

1.

The problem is equivalent to minimize $\|x - z\|$ and the function $x \mapsto \|x - z\|$ is continuous, so it has a minimum on every non-empty compact set (Weierstrass. For a recall of the argument, see *Part 2. Question 1.*). By supposition, Q is non empty, so there exists a $R > 0$ such that $B(z, R) \cap Q$ is non empty and contain a certain q (where $B(z, R)$ is the closed ball of radius R centered in z). Note that now if $x \notin B(z, R)$, then $\|x - z\| > R \geq \|q - z\|$ and x is not the min because q is a better minimizer. We can reduce the problem in $\tilde{Q} := B(z, R) \cap Q$:

$$\min_{x \in Q} \|x - z\| = \min_{x \in Q \cap B(z,R)} \|x - z\|.$$

Since Q and $B(z, R)$ are closed, $B(z, R) \cap Q$ is closed and bounded by R, it is compact. By construction it is non-empty and as a result, it as a minimum : $\text{Proj}_Q(z)$ is non-empty.

2.

– Suppose $\text{Proj}_C(z) = \{0\}$. Then 0 is a global and in particular local minimum of

$$\min_{x \in C} \frac{1}{2}\|x - z\|^2$$

Since f is diffenrentiable, we know by theorem 7.18 of the course that 0 is a stationary point and we obtain

$$z = -(0 - z) = -\nabla f(0) \in (T_0 C)^\circ = C^\circ,$$

because we can show that actually $T_0 C = C$. Indeed, for every point x in C, we can take the curve $c(t) = tx \in C \quad \forall t \geq 0$ (C is a cone), that converge to 0 when $t \to 0^+$ and such that $c'(0) = x$, so $x \in T_0 C$ and $C \subset T_0 C$. Alternatively let's consider a point x in $T_0 C$ with associated sequence $\frac{x_n - 0}{t_n} \to x$. By definition of a tangent cone, $x_n \in C$, and by definition of a cone, $x_n / t_n \in C$. As a result we have that the sequence must converge inside C because it is closed. We have the second inclusion $C \supset T_0 C$ and conclude that equality holds.

– Suppose $z \in C^\circ$, then for all non-null point x in C,

$$\|x - z\|^2 = \underbrace{\|x\|^2}_{>0} - 2 \underbrace{\langle x, z \rangle}_{\leq 0} + \|z\|^2 > \|z\|^2 = \|0 - z\|^2$$

and 0 is a strictly better minimizer than any other point x. We conclude that

$$\text{Proj}_C(z) = \arg\min_{z\in C} \frac{1}{2}\|x-z\|^2 = \{0\}.$$

These two parts prove the result.

3.

Using question 2, since $T_{x^*}S$ is a non-empty closed cone

$$\text{Proj}_{T_{x^*}S}(-\nabla f(x^*)) = \{0\} \iff -\nabla f(x^*) \in (T_{x^*})^\circ = N_{x^*},$$

and that is x^* being stationary by definition.

4.

a) If $v \in \text{Proj}_C(z)$, then it is by definition a minimizer of $\min_{x\in C} 1/2\|x-z\|^2$ and by theorem 7.18, a stationary point because $x \mapsto 1/2\|x-z\|^2$ is differentiable with gradient $v-z$ (evaluated in v). By definition, we have $-(v-z) \in N_vC$, i.e. for all $d \in T_vC$, $\langle d, z-v\rangle \leq 0$. But clearly $\{v, -v\} \in T_vC$ because of the tangent curves $c_{1,2}(t) = v \pm tv = (1 \pm t)v$ which are in C as soon as $t < 1$. So taking both cases $d = \pm v$,

$$0 \geq \langle \pm v, z-v\rangle = \pm\langle v, z-v\rangle$$

and $\langle v, z-v\rangle$ must be 0.

b) For an arbitrary projection v, we have

$$\|v\|^2 = \|(v-z)+z\|^2 = \|v-z\|^2 + 2\underbrace{\langle v-z, z\rangle}_{=0 \text{ by a)}} + \|z\|^2 = \min_{x\in C}\|x-z\|^2 + \|z\|^2$$

doesn't depend on v.

5.

We take $f = \text{Id}$ and $S = \mathbb{R}_+$. We have $-\nabla f = -1$ and

$$T_xS = \begin{cases} \mathbb{R}_+ \text{ if } x = 0 \\ \mathbb{R} \text{ if } x > 0 \end{cases}$$

We get

$$\text{Proj}_{T_0}(-\nabla f(0)) = \text{Proj}_{\mathbb{R}_+}(-1) = 0$$
$$\text{Proj}_{T_x}(-\nabla f(x)) = \text{Proj}_{\mathbb{R}}(-1) = -1 \quad \forall x > 0.$$

Its norm is clearly not continuous in 0.

6.

We compute an explicit form of the projection.

First of all, the set S is one dimensional and can be described as $h^{-1}(\{0\})$ with the single equality constraint $h(z) = \|z\|^2 - 1$. The tangent cone is the tangent line to the circle , because all associated tangent sequences follow the same line. More precisely, this is follow from theorem 8.14 about linear independence constraint qualification (LICQ), since here we have only one constraint function, the constraint gradient make a trivial independent collection. Thus a qualification condition holds and $T_zS = F_zS = (\nabla h(z))^\perp = (2z)^\perp$.

A non-null perpendicular vector of $z = (x, y) \in S$ is $(y, -x) =: i(z)$, and the the tangent cone is then $\mathbb{R}i(z)$. We know by Gram-Schmidt that the projection of a vector a to a vector b is just $\bar{a} = \frac{\langle a,b\rangle}{\|b\|^2}b$. In our case, $\|z\| = 1 = \|i(z)\|$ and

$$\mathrm{Proj}_{T_z}(-\nabla f(z)) = \mathrm{Proj}_{\mathbb{R}i(z)}(-\nabla f(z)) = -\frac{\langle \nabla f(z), i(z)\rangle}{\|i(z)\|^2} i(z) = -\langle \nabla f(z), i(z)\rangle i(z)$$

and its norm is $q(z) = |\langle \nabla f(z), i(z)\rangle|$. Since i , ∇f, and the scalar product are all continuous, q is continuous.

7.

a) The fact that LICQ holds, means that the jacobian $J(x) := Dh(x)$ has always full row rank p in S. By theorem 8.14, we have

$$T_xS = F_xS = \{y \in \mathbb{R}^n | \langle y, \nabla h_i(x)\rangle = 0 \quad \forall i = 1, \dots, p\} = \ker(J(x)).$$

We will need later to compute projection on this null space, so we use a matrix representation for it, using the singular decomposition and the pseudo inverse.

We recall that all matrix $A \in \mathbb{R}^{m\times n}$ can be written $A = PDQ$ where $P \in \mathbb{R}^{m\times m}$ and $Q \in \mathbb{R}^{n\times n}$ are orthogonal, and $D \in \mathbb{R}^{m\times n}$ is diagonal. The matrix D^+ is defined with $(D^+)_{ij} = (D_{ij})^{-1}$ when $(D^+)_{ij}$ is non null, and 0 otherwise. The pseudo inverse matrix A^+ is then PD^+Q and has the properties that $AA^+A = A$ and A^+A is symetric. Moreover when A has full row rank, the pseudo inverse $A^+ = A^\top(AA^\top)^{-1}$.

Now we proove that $\ker(A) = \mathrm{im}(I - A^+A)$. Indeed, if x is in the null space, then it can be written

$$x = x - A^+0 = x - A^+Ax = (I - A^+A)x$$

and $x \in \mathrm{im}(I - A^+A)$.
Alternatively, if $y = (I - A^+A)x$, then

$$A(I - A^+A)x = (A - AA^+A)x = (A - A)x = 0$$

and $y \in \mathrm{im}(I - A^+A)$.

Finally we get $T_xS = \mathrm{im}(I - A^+A)$.

b) Thanks to a), we can rewrite the minimisation problem $\mathrm{Proj}_{T_xS}(z)$ as follow:

$$\min_{v\in T_xS} \|v - z\| = \min_{v\in \mathrm{im}(I-A^+A)} \|v - z\| = \min_{y\in\mathbb{R}^n} \|(I - J(x)^+J(x))y - z\|$$

By the last squares, we know that the term $v = (I - J(x)^+J(x))y$ is the same for all minimizers y, and a solution is given by the equation

$$(I - J(x)^+J(x))^\top(I - J(x)^+J(x))y = (I - J(x)^+J(x))^\top z \tag{1}$$

But now we can show that actually

$$(I - J(x)^+ J(x))^\top (I - J(x)^+ J(x)) = (I - J(x)^+ J(x))$$

Indeed, $I - J(x)^+ J(x)$ is symetric, and we have

$$\begin{aligned}(I - J(x)^+ J(x))^\top (I - J(x)^+ J(x)) &= (I - J(x)^+ J(x))^2 = I - 2J(x)^+ J(x) + J(x)^+ J(x) J(x)^+ J(x) \\ &= I - 2J(x)^+ J(x) + J(x)^+ J(x) = I - J(x)^+ J(x).\end{aligned}$$

As a result, the equation (1) read actually $v = (I - J(x)^+ J(x))y = (I - J(x)^+ J(x))z$ and we have a direct explicit expression of the projection:

$$\text{Proj}_{T_x S}(z) = (I - J(x)^+ J(x))z$$

c) Since J and ∇f are continuous, and looking at the expression $-(I - J(x)^+ J(x))\nabla f(x)$, $\text{Proj}_{T_x S}(-\nabla f(x))$ is continuous in x if the map $A \mapsto A^+$ is continuous too. It is not the case in general, but in our case $A = J(x)$ has always full row rank and a linear algebra theorem tell us it can be written $A^+ = A^\top (AA^\top)^{-1}$, which is continuous in A(matrix sum, product, transpose, inverse are continuous). As the result the projection itself is actually continuous and

$$\text{Proj}_{T_x S}(z) = -(I - J(x)^\top (J(x) J(x)^\top)^{-1} J(x)) \nabla f(x)$$

with norm of course continuous too.

Part 2 : A Frank–Wolfe algorithm

1.

The problem (2) always has a solution. First, S is non empty. Second, define $g : \mathcal{E} \to \mathbb{R}$, by $g(x) = \langle w, x \rangle$. The function g is continuous since the scalar product is. Then we know there exists an infimum of the function g, in other words there exists a converging sequence $(x_n)_{n=1}^{\infty} \subset S$ such that

$$\lim_{n\to\infty} \langle w, x_n \rangle = \inf_{x \in S} \langle w, x \rangle.$$

As S is compact, $(x_n)_{n=1}^{\infty}$ must converge in S. Therefore the sequence does not only attain the infimum, but the mininum. To conclude, the problem (2) always has a solution.

2.

By taking $f : \mathbb{R}^2 \to \mathbb{R}$, defined by $f(x, y) = x$ (f is continuously differentiable, is convex and ∇f is lipschitz continuous), we have $w = \nabla f(x, y) = (1, 0)^\top$ for all $(x, y) \in \mathbb{R}^2$. Take $S = \{(x, y) \in \mathbb{R}^2 : \max(|x|, |y|) \leq 1\}$ (the square in $\mathbb{R}^2$). This set is indeed convex and compact, since it is bounded and closed.

We have, for all (x, y) in $\mathbb{R}^2$:

$$\langle w, (x, y) \rangle = x,$$

therefore problem (2) can be rewritten as

$$\min_{(x,y)\in S} x.$$

Clearly the solutions are given by $\{(-1, y) : |y| \leq 1\}$. The problem admits more than one solution.

3.

It is important because the convexity of S assures $x_{k+1} = (1-\eta_k)x_k + \eta_k s(x_k)$ to be in S only if $\eta_k \in [0, 1]$, as we know that x_k and $s(x_k)$ are in S by construction.

4.

(B1): as $f : \mathcal{E} \to \mathbb{R}$ is continuously differentiable and ∇f is L-lipschitz continuous (see theorem 3.2 in the notes, with $u = x_{k+1} - x_k$ and $x = x_k$).
(B2): by definition of

$$x_{k+1} = x_k - \eta_k x_k + \eta_k s(x_k) \iff x_{k+1} - x_k = \eta_k (s(x_k) - x_k)$$

Therefore:

$$\|x_{k+1} - x_k\|^2 \leq |\eta_k|^2 \|s(x_k) - x_k\|^2 \leq |\eta_k|^2 d_S^2.$$

The last inequality is implied by definition of the diameter of S d_S, as

$$\|s(x_k) - x_k\| \leq \max_{x,y\in S} \|x - y\| = d_S.$$

(B3): by definition of $s(x_k) = \text{argmin}_{x\in\Delta^n} \langle \nabla f(x_k), x \rangle$, we have

$$\langle \nabla f(x_k), s(x_k) \rangle \leq \langle \nabla f(x_k), x^* \rangle$$

which is equivalent to

$$\langle \nabla f(x_k), s(x_k) - x_k \rangle \leq \langle \nabla f(x_k), x^* - x_k \rangle$$

by linearity of the scalar product. As η_k is in between zero and one, we have the expected inequality as:

$$\nabla f(x_k)^\top (s(x_k) - x_k) \leq \nabla f(x_k)^\top (x^* - x_k).$$

(B4): by theorem 4.21, as $f : \mathcal{E} \to \mathbb{R}$ is differentiable on a euclidean space $\mathcal{E}$ and convex, we have :

$$\forall x, y \in \mathcal{E} : f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle.$$

In particular for $x = x_k$ and $y = x^*$:

$$\langle \nabla f(x_k), x^* - x_k \rangle \leq f(x^*) - f(x_k) \iff \nabla f(x_k)^\top (x^* - x_k) \leq f(x^*) - f(x_k)$$

which is the expected inequality.

5.

We have:

$$f(x_1) - f(x^*) \leq f(x_1) - f(x_0) - \big(f(x^*) - f(x_0)\big).$$

By using the previous exercise on $f(x_1) - f(x_0)$:

$$f(x_1) - f(x^*) \leq \eta_0(f(x^*) - f(x_0)) + \frac{L}{2}\eta_0^2 d_S^2 - \big(f(x^*) - f(x_0)\big) = \frac{L}{2} d_S^2,$$

the last equality being given as $\eta_0 = 1$.

6.

We show this result by induction. By 5., the result is true for $k = 1$. Now for $k \geq 1$:

$$f(x_{k+1}) - f(x^*) = f(x_{k+1}) - f(x_k) - \big(f(x^*) - f(x_k)\big)$$

By using 4. on $f(x_{k+1}) - f(x_k)$:

$$f(x_{k+1}) - f(x^*) \leq \eta_k\big(f(x^*) - f(x_k)\big) + \frac{L}{2}\eta_k^2 d_S^2 - \big(f(x^*) - f(x_k)\big) = (\eta_k - 1)\big(f(x^*) - f(x_k)\big) + \frac{L}{2}\eta_k^2 d_S^2$$

Also, we know that by definition of η_k:

$$\eta_k - 1 = \frac{2}{k+2} - 1 = \frac{2 - (k+2)}{k+2} = \frac{-k}{k+2}.$$

Therefore:

$$f(x_{k+1}) - f(x^*) \leq \frac{k}{k+2}\big(f(x_k) - f(x^*)\big) + \frac{L}{2}\eta_k^2 d_S^2 \leq \frac{k}{k+2}\frac{2Ld_S^2}{k+2} + \frac{L}{2}\Big(\frac{2}{k+2}\Big)^2 d_S^2 = \Big(\frac{2k}{(k+2)^2} + \frac{2}{(k+2)^2}\Big) L d_S^2,$$

where we used the induction hypothesis on the second inequality.

Then, we have for all $k \geq 1$:

$$6 \leq 8 \iff 2k^2 + 6k + 2k + 6 \leq 2k^2 + 8k + 8 \iff (2k+2)(k+3) \leq 2(k+2)^2 \iff \frac{2k+2}{(k+2)^2} \leq \frac{2}{k+3}.$$

Hence we can conclude:

$$f(x_{k+1}) - f(x^*) \leq \frac{2k+2}{(k+2)^2} L d_S^2 \leq \frac{2}{k+3} L d_S^2 = \frac{2}{(k+1)+2} L d_S^2$$

which is the expected result.

7.

Δ^n is **convex** by Corollary 9.22 of the course. Indeed, define:

- $g_i : \mathcal{E} \to \mathbb{R}$ for $i = 1, ..., n$ by $g_i(\vec{x}) = -x_i$. They are convex as they are linear.
- $h : \mathcal{E} \to \mathbb{R}$ by $h(\vec{x}) = \sum_{i=1}^n x_i - 1 = \vec{1}^\top \vec{x} - 1$ which is affine.

Then the set $\Delta^n = \{\vec{x} \in \mathcal{E} : h(\vec{x}) = 0, g_i(\vec{x}) \leq 0 \forall i = 1, ..., n\}$ is convex.

Δ^n is **compact** by Heine-Borel theorem, as $\Delta^n \subset \mathcal{E}$ an euclidean space. Δ^n is compact if and only if it is closed and bounded.

It is clearly bounded as for $x \in \Delta^n$, then all components of x are in between 0 and 1.

It is closed sequentially (therefore closed). Take $(x_k)_{k \geq 1} \subset \Delta^n$ such that $\lim_{k \to \infty} x_k = x$ for $x \in \mathcal{E}$. As g_i are continuous for all $i = 1, ..., n$ we have $g_i(x) = \lim_{k \to \infty} g_i(x_k) \leq 0$. In the same way, as h is continuous, we have $h(x) = \lim_{k \to \infty} h(x_k) = 0$. Therefore x belongs to Δ^n and Δ^n is closed.

Δ^n is **non-empty**, as $\vec{x} = (1, 0, ...0)^\top \in \Delta^n$.

8.

We look for

$$\min_{x \in \Delta^n} \langle w, x \rangle = \min_{1 \leq i \leq n} w_i$$

as $x_i \geq 0$ for all $i = 1, ..., n$, and that x gives weight for each component of w. We can show it more precisely by contradiction. Suppose there exists $\hat{x} \in \Delta^n$ such that

$$\langle w, \hat{x} \rangle = \min_{x \in \Delta^n} \langle w, x \rangle < \min_{1 \leq i \leq n} w_i$$

Let $m = \operatorname{argmin}_{1 \leq i \leq n} w_i$. We have, by definition of the argmin, $w_m \leq w_i \forall i = 1, ..., n$. Then we can rewrite the strict inequality above as:

$$\sum_{i=1}^n w_i \hat{x}_i < w_m$$

and we have the contradiction as

$$w_m \sum_{i=1}^n \hat{x}_i \leq \sum_{i=1}^n w_i \hat{x}_i < w_m$$

which would imply

$$\sum_{i=1}^n \hat{x}_i < 1.$$

To attain the smallest value, we can take $(x_i)_{i=1}^n = \delta_{im}$, with $m = \operatorname{argmin}_{1 \leq i \leq n} w_i$.

9.

With $S = \Delta^n$, by 8., it suffices to solve $\min_{1 \leq i \leq n} w_i$ The computational complexity of computing my solution is $O(n)$, as the computer needs to go through all components of w once. To show this, I implemented `min_complexity.m`, which measures the time that the computer needs in order to take the index of the minimum component of a vector. In the code below, we take a random vector from which we vary the size. The length of the vector is multiplied by two at each loop. We obtain the plot below, which is in `loglog`. We observe that the slope of the time is similar to the slope of the linear function, therefore the function `min` runs in $O(n)$.

Note that we run the `min` function 15 times in order to calculate a mean of the time the function takes (see lines 12 to 18 in the code below). The number of times is arbitrary.

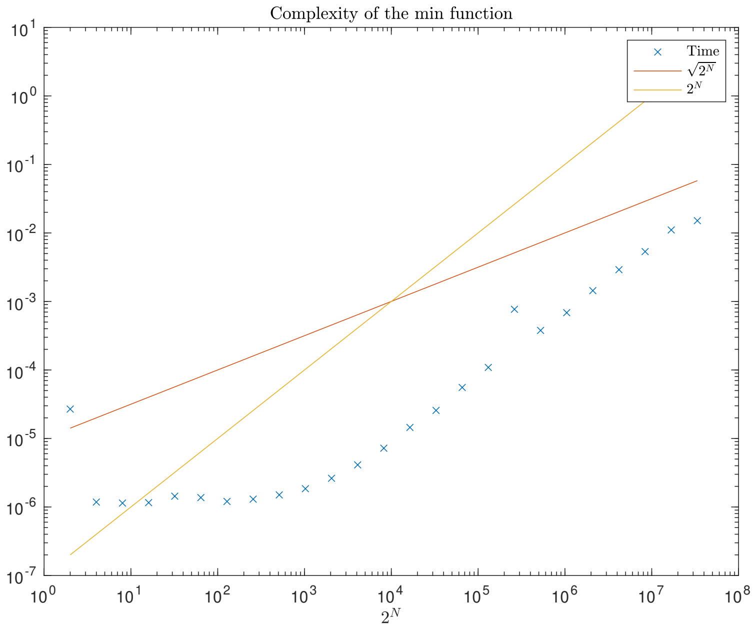

Figure 1: Graph of exercise 9

```
clc
clear
close

N_max=25;
N=1:1:N_max;
time=zeros(size(N,1),1);

for n=N
    m=2^n;
    t=0;
    for j=1:15
        x=randn(m,1);
        tic
        [~,I]=min(x);
        t=t+toc;
    end
    time(n)=t/15;
end
%% plot
loglog(2.^N,time,'x',2.^N,1e-5*sqrt(2.^N),2.^N,1e-7*(2.^N))
legend('Time','$\sqrt{2^N}$','$2^N$','Interpreter','latex')
title('Complexity of the min function','Interpreter','latex');
xlabel('$2^N$','Interpreter','latex');
```

10.

We have that $f(x)$ is continuous by the continuity of the norm, the square and the matrix multiplication. Since Δ^n is convex, compact and non-empty, by the same reasoning as in 1., this problem admits a solution. Indeed, we can again define a sequence $(x_n)_{n\geq 1} \subset \Delta^n$ such that

$$\lim_{n\to\infty} f(x_n) = \inf_{x\in\Delta^n} f(x).$$

As Δ^n is compact, the sequence $(x_n)_{n\geq 1}$ must converge in Δ^n, therefore the minimum is attained in Δ^n.

Moreover, the solution might not be unique. Suppose A is the null matrix and b is the null vector. Then all $x \in \Delta^n$ are solutions of the minimization problem as $Ax - b = 0$ for all $x \in \Delta^n$.

11.

We have :

$$f(x) = 1/2(\langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle)$$

by properties of the scalar product, in particular the symmetry of it. Then take $x, v \in \Delta^n$, and $t > 0$. Let's calculate:

$$f(x+tv) - f(x) = 1/2\big(\langle A(x+tv), A(x+tv)\rangle - 2\langle A(x+tv), b\rangle + \langle b, b\rangle - \langle Ax, Ax\rangle + 2\langle Ax, b\rangle - \langle b, b\rangle\big)$$

$$= 1/2\big(\langle Ax, Ax\rangle + 2t\langle Ax, Av\rangle + t^2\langle Av, Av\rangle - 2\langle Ax, b\rangle - 2t\langle Av, b\rangle - \langle Ax, Ax\rangle + 2\langle Ax, b\rangle\big)$$

$$= 1/2\big(2t\langle Ax, Av\rangle + t^2\langle Av, Av\rangle - 2t\langle Av, b\rangle\big) = t\langle Ax, Av\rangle + \frac{1}{2}t^2\langle Av, Av\rangle - t\langle Av, b\rangle$$

Therefore:

$$Df(x)[v] = \lim_{t\to 0} \frac{f(x+tv) - f(x)}{t} = \lim_{t\to 0}\langle Ax, Av\rangle + \frac{1}{2}t\langle Av, Av\rangle - \langle Av, b\rangle$$

$$= \langle Ax, Av\rangle - \langle b, Av\rangle = \langle Ax - b, Av\rangle = \langle A^\top(Ax - b), v\rangle$$

by symmetry and linearity of the scalar product. To conclude, since $Df(x)[v] = \langle \nabla f(x), v\rangle$, we have

$$\nabla f(x) = A^\top(Ax - b).$$

12.

We have $g : [0, 1] \to \mathbb{R}$. First, let's rewrite $g(\eta)$, for all x and y in Δ^n:

$$g(\eta) = \frac{1}{2}\left\|A[(1-\eta)x + \eta y] - b\right\|^2 = \frac{1}{2}\left\|Ax - b + \eta A(y - x)\right\|^2$$

$$= \frac{1}{2}\langle \eta A(y-x) + Ax - b, \eta A(y-x) + Ax - b\rangle$$

$$= \frac{1}{2}\big(\eta^2\langle A(y-x), A(y-x)\rangle + 2\eta\langle A(y-x), Ax - b\rangle + \langle Ax - b, Ax - b\rangle\big).$$

In short:

$$g(\eta) = 1/2\big[\eta^2\left\|A(y-x)\right\|^2 + 2\eta\langle A(y-x), Ax - b\rangle + \left\|Ax - b\right\|^2\big].$$

Therefore, $g(\eta)$ is a quadratic function of η. Moreover, as $\frac{1}{2}\left\|A(y-x)\right\|^2 \geq 0$, it is convex. Finding the minimum of $g(\eta)$ corresponds to find the minimum of f on the segment defined by x and y, which is what we are looking for. The minimum is either on the boundaries of $[0, 1]$ or in the interior. If it was on the interior, the minimum of g would be given by the zero of the gradient of g:

$$\delta_\eta g(\eta) = \eta\langle A(y-x), A(y-x)\rangle + \langle A(y-x), Ax-b\rangle.$$

The optimal value would be given by :

$$\bar{\eta} = -\frac{\langle A(y-x), Ax-b\rangle}{\|A(y-x)\|^2}$$

If the minimum was on the boundary, the optimal value would be either zero or one.
The minimum is attained at $\bar{\eta}$, at zero, or at one. We have the following three cases:

- $\bar{\eta} \leq 0$: then by representing the quadratic function with positive coefficient in front of the quadratic term, we can easily see that the minimum is attained at zero. See below an illustration of a quadratic function with positive coefficient in front of the square term. We observe that $\bar{\eta}$ is equal to -2, and that the minimum is indeed attained at zero in the interval $[0, 1]$.

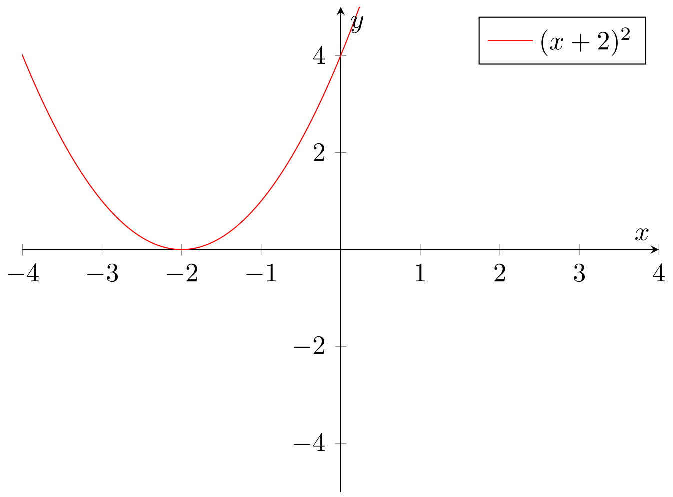

- $0 < \bar{\eta} < 1$: the value $\bar{\eta}$ is in the interior of $[0, 1]$. Therefore the minimum is attained at $\bar{\eta}$. See an illustration below, with $\bar{\eta} = 1/2$.

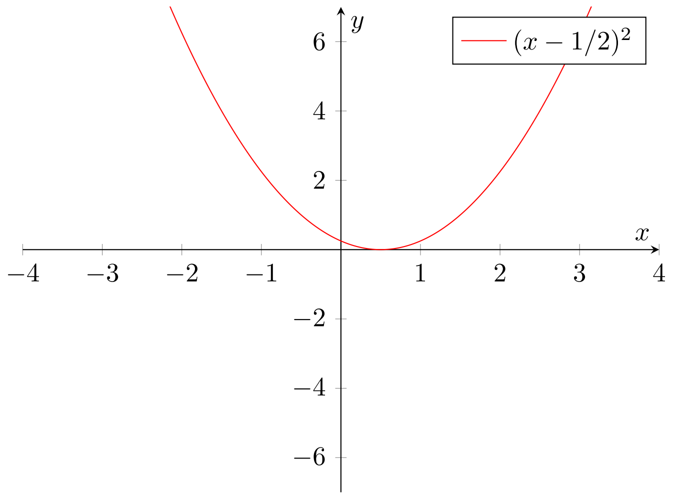

- $\bar{\eta} \geq 1$: then by representing the quadratic function with positive coefficient in front of the quadratic term, we can easily see that the minimum is attained at one. See below for illustration.

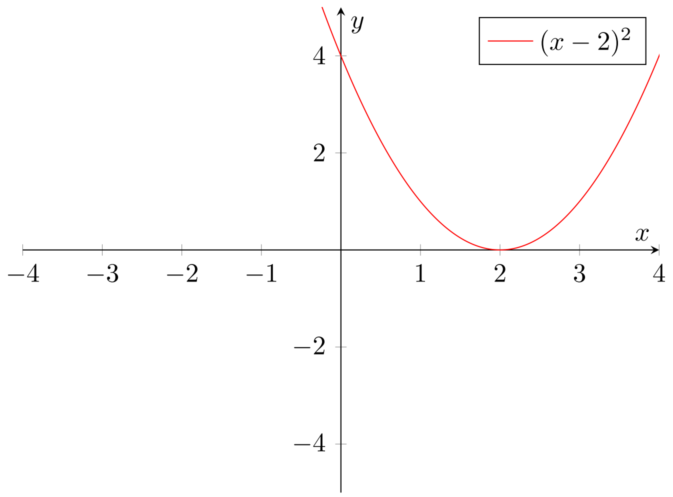

13.

Here is the algorithm:

Algorithm 1 Frank-Wolfe

Input: $S = \Delta^n$, $f(x) = \frac{1}{2}\|Ax - b\|^2$, $x_0 = (1, 0, ...0) \in \mathbb{R}^n$
for k=0,1,... **do**
  Compute $j = \text{argmin}_{i=1,...,n} \nabla f(x_k)$
  $(s(x_k))_{i=1}^n = \delta_{ij}$
  $\eta_k = \max(0, \min(-\frac{\langle A(s(x_k)-x_k), Ax_k-b\rangle}{\|A(s(x_k)-x_k)\|^2}, 1))$
  $x_{k+1} = (1-\eta_k)x_k + \eta_k s(x_k)$
end for

14.

We implement the function `FrankWolfe.m` below:

```
function [Gap,k,x_k]=FrankWolfe(maxit,tol,A,b,n)
% run the Frank-Wolfe algorithm for a function f(x)=1/2||Ax-b||^2 , x in
% R^n. It returns a vector Gap (the F.-W. gap at each iteration), the
% maximum of iteration k the algorithm has reached, and the solution x_k
% obtained at the end of the algorithm.

f=@(x) 0.5*dot(A*x-b,A*x-b);
grad=@(x) A'*(A*x-b);
x_k=zeros(n,1);
x_k(1)=1;
Gap=zeros(maxit,1);

for k=1:maxit
    grad_k=grad(x_k);
    [~,index]=min(grad_k);
    sx_k=zeros(n,1);
    sx_k(index)=1;
    nu_k=-dot(A*(sx_k-x_k),A*x_k-b)/dot(A*(sx_k-x_k),A*(sx_k-x_k));
    nu_k=max(0,min(nu_k,1));
    x_k=(1-nu_k)*x_k+nu_k*sx_k;
    %Frank-Wolfe gap:
    g_k=dot(grad_k,x_k-sx_k);
    Gap(k)=g_k;
    if g_k<tol
        break
    end
end
end
```

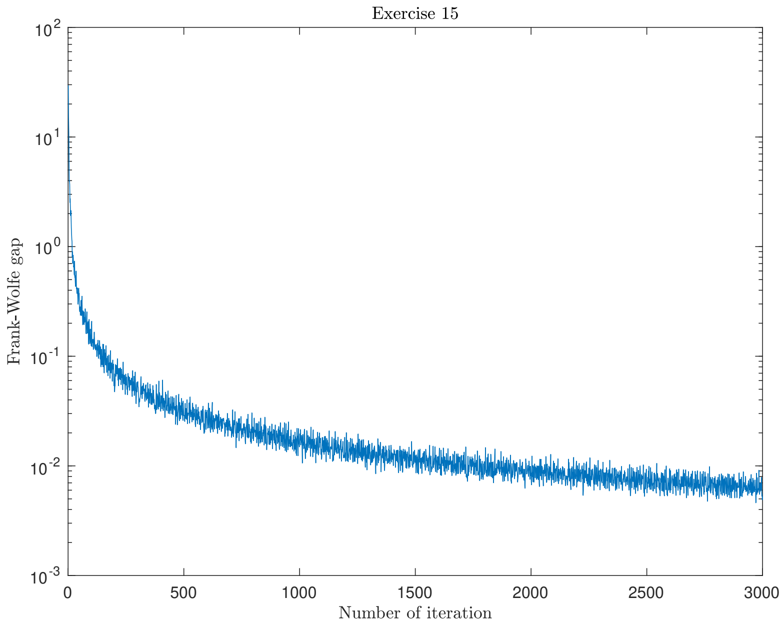

Figure 2: Graph of exercise 15

15.

We implement `ex15.m` which runs the Frank-Wolfe algorithm on the given data.

```
clc
clear
close
% exercise 15.

load data.mat
maxit=3000;
tol=1e-3;

[Gap,k,~]=FrankWolfe(maxit,tol,A,b,n);

semilogy(1:1:k,Gap(1:k))
title('Exercise 15','Interpreter','latex')
xlabel('Number of iteration','Interpreter','latex');
ylabel('Frank-Wolfe gap','Interpreter','latex');
```

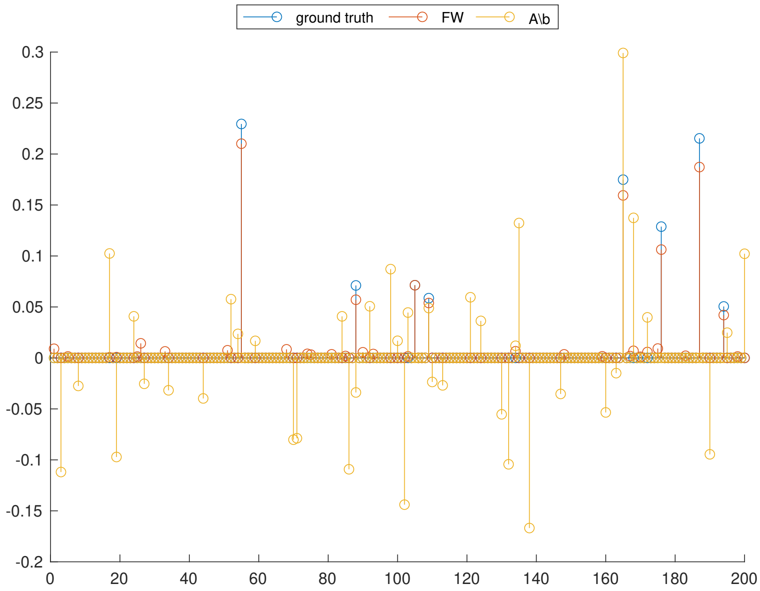

Figure 3: Graph of exercise 16

16.

We implement `ex16.m`, see below. We obtain the plot given by the function `plot_data.m`.

```
clc
clear
close
% exercise 16.

load data.mat
maxit=3000;
tol=1e-3;

%solution given by FrankWolfe:
[~,~,sol_FW]=FrankWolfe(maxit,tol,A,b,n);
nnz(sol_FW)

%solution given by solving the linear system:
sol_lin=A\b;
nnz(sol_lin)


plot_data(x,sol_FW,sol_lin)
```

We observe on the graph that the solution given by the backslash command gives worse result then the solution given by the Frank-Wolfe method. Indeed, observe the yellow branches going down as no blue branches have this behaviour. Notice the red branches being closer to the blue branches.

Also using the `nnz` function, which gives the number of non-zero elements in a vector, we note that the solution given by Frank-Wolfe has 34 of them, and the one given by solving the linear system has 40 of non-zero entries (in fact x has 8 non-zeros entries). So the Frank-Wolfe solution is more sparse than the other one.

To conclude, the Frank-Wolfe algorithm gives a better solution than the one given by the backslash command.

Part 3 : KKT conditions and constraint qualifications.

In this part, we set $F = (f_1, \dots, f_N)^\top$ and $\mathbb{1} = (1, \dots, 1)^\top \in \mathbb{R}^N$, giving the descriptions $f = \|F\|_\infty$ and $S = \{(x, y) \in \mathbb{R}^{n+1} | F(x) \leq y\mathbb{1}\}$.

1.

For all $(x, y) \in S$, $F(x) \leq y\mathbb{1}$ and so $f(x) \leq y$. As a result

$$\min f \leq \min_{(x,y)\in S} y$$

Furthermore, for all minimizer x of f, we have trivially that $(x, f(x)) \in S$ because $F(x) \leq \|F(x)\|_\infty \mathbb{1}$, so

$$\min_{(x,y)\in S} y \leq f(x) = \min f.$$

the two inequalities give us the result.

2.

We set the function $g(x, y) = F(x) - y\mathbb{1}$, with gradients $\nabla g_i(x, y) = \begin{pmatrix} \nabla f_i(x) \\ -1 \end{pmatrix}$ and $\nabla_{xy}(y) = \begin{pmatrix} (0) \\ -1 \end{pmatrix}$.

The KKT conditions are $\begin{cases} -\begin{pmatrix} (0) \\ 1 \end{pmatrix} = \sum_{i=1}^N \lambda_i \begin{pmatrix} \nabla f_i(x) \\ -1 \end{pmatrix} \\ \lambda_i(f_i(x) - y) = 0 \quad \forall i \in \{1, \dots, N\} \end{cases}$ i.e. $\begin{cases} \sum_{i=1}^N \lambda_i \nabla f_i(x) = 0 \\ \sum_{i=1}^N \lambda_i = 1 \\ \lambda_i(f_i(x) - y) = 0 \quad \forall i \in \{1, \dots, N\} \end{cases}$

For some $\lambda_i \geq 0$.

3.

We choose $\begin{cases} f_1(x) = 2x_1 \\ f_2(x) = 2x_2 \\ f_3(x) = x_1 + x_2 \end{cases}$, and we get $\begin{cases} \nabla g_1(x, f(x)) = (2, (0), -1))^\top \\ \nabla g_2(x, f(x)) = (0, 2, (0), -1))^\top \\ \nabla g_3(x, f(x)) = (1, 1, (0), -1))^\top \end{cases}$.

Now $\nabla g_1 + \nabla g_2 - 2\nabla f_3 = 0$ everywhere, and in particular in the active points $(x, f(x))$, like ((0),0) for example.

4.

To have MFCQ, since here we don't have the equalities constraints, we just have to look to the inequality constraint gradients.
Let $z = (x, y) \in S$ and $I(x, y) = \{j \in \{1, \dots, N\} | f_i(x) = y\}$. Then for all $i \in I(z)$, $\nabla g_i(z) = (\nabla f_i(x)^\top, -1)^\top$ and $((0), 1)\nabla g_i(z) = -1 < 0$. The MFCQ holds thanks to the direction $((0), 1)$.

5.

By question 4., we have a constraint qualification that holds, MFCQ. The KKT Theorem tell us that (x^*, y^*) is a KKT point. Looking at question 2, we see that it means that 0 is in the convex envelop of the active gradients $\nabla f_i(x^*)$, with the λ's being the coordinates of 0 for such a convex linear combination. In particular they are linearly dependent and the Lagrange multipliers might not be unique. More precisely, continuously infinitely possible multipliers can exist.