EE127A 3/27/09
L. El Ghaoui
YOUR NAME HERE: **SOLUTIONS**
YOUR SID HERE: **42**

Midterm 1 Solutions

The exam is open notes, but access to the Internet is not allowed. The maximum grade is 20. When asked to prove something, do not merely take an example; provide a rigorous proof with clear steps. Also, some parts are more difficult than others and you are not expected to finish the exam.

1. (*2 points*) Show that the Frobenius norm of a matrix A depends only on its singular values. Precisely, show that

$$\|A\|_F = \|\sigma\|_2,$$

where $\sigma := (\sigma_1, \ldots, \sigma_r)$ is the vector formed with the singular values of A, and r is the rank of A.

Solution: We present two solutions. Assume $A \in \mathbf{R}^{m\times n}$.

a) By re-writing this in terms of the squared Frobenius norm,

$$\|A\|_F^2 = \mathrm{tr}(A^TA) = \sum_{i=1}^{n} \lambda_i(A^TA) = \sum_{i=1}^{r} \lambda_i(A^TA) = \sum_{i=1}^{r} \sigma_i^2 = \|\sigma\|_2^2\,,$$

the sum going to r because only the first r eigenvalues of A are non-zero.

b) We can use the SVD of $A = U\tilde{\Sigma}V^T$, with $\tilde{\Sigma} = \mathbf{diag}(\sigma_1, \ldots, \sigma_r, 0, \ldots, 0)$, directly. We have

$$\|A\|_F^2 = \|U\tilde{\Sigma}V^T\|_F^2 = \mathrm{tr}(V\tilde{\Sigma}U^TU\tilde{\Sigma}V^T) = \mathrm{tr}(V\tilde{\Sigma}^2V^T) = \mathrm{tr}(\tilde{\Sigma}^2V^TV) = \mathrm{tr}\,\tilde{\Sigma}^2,$$

which proves the result.

2. (*8 points*) We are given a matrix $A \in \mathbf{R}^{m\times n}$. We consider a matrix least-squares problem

$$\min_X \; \|AX - I_m\|_F\,, \tag{1}$$

where the variable is $X \in \mathbf{R}^{n\times m}$, and I_m is the $m \times m$ identity matrix.

   (a) Show that the problem can be reduced to a number of ordinary least-squares problems, each involving one column of the matrix X. Make sure you define precisely the form of these least-squares problems.

   Solution: Let $x_j \in \mathbf{R}^n$ denote the j-th column of A, so that $X = [x_1, \ldots, x_m]$. We have $AX - I_m = [Ax_1 - e_1, \ldots, Ax_m - e_m]$, where e_i is the i-th basis vector in

$\mathbf{R}^m$. Since the squared Frobenius norm of a matrix is the sum of squares of the Euclidean norms of its columns, the objective function of our problem is

$$\|AX - I_m\|_F^2 = \|[Ax_1 - e_1, \ldots, Ax_m - e_m]\|_F^2 = \sum_{i=1}^m \|Ax_i - e_i\|_2^2.$$

Note that the vector x_i appears only once, in the i-th term of the sum. Hence, to minimize the above objective function, we can simply minimize the i-th term with respect to x_i, *independently of the other terms.* That is:

$$\min_{X=[x_1,\ldots,x_m]} \|AX - I_m\|_F^2 = \min_{x_1,\ldots,x_m} \sum_{i=1}^m \|Ax - e_i\|_2^2 = \sum_{i=1}^m \min_x \|Ax - e_i\|_2^2.$$

Hence, to solve our problem we can simply solve m ordinary least-squares problems, of the form

$$\min_x \|Ax - e_i\|_2, \quad i = 1, \ldots, m.$$

If x_i^* is an optimal solution for the above, then $X = [x_1^*, \ldots, x_m^*]$ is optimal for our original problem.

(b) Show that, when A is full column rank, then the optimal solution is unique, and given by

$$X^* = (A^T A)^{-1} A^T.$$

Solution: When A is full rank, the solution to the least-squares problem

$$\min_x \|Ax - y\|_2$$

is unique, and given by $x^* = (A^T A)^{-1} A^T y$.
Applying this to $y = e_i$, $i = 1, \ldots, m$, we obtain that the optimal X is unique, and given by $X^* = [x_1^*, \ldots, x_m^*]$, with $x_i^* = (A^T A)^{-1} A^T e_i$, $i = 1, \ldots, m$. That is:

$$\begin{aligned} X^* &= \left[(A^T A)^{-1} A^T e_1 \quad \cdots (A^T A)^{-1} A^T e_m\right] \\ &= (A^T A)^{-1} A^T \cdot \left[e_1 \quad \cdots e_m\right] \\ &= (A^T A)^{-1} A^T \cdot I_m = (A^T A)^{-1} A^T. \end{aligned}$$

Another method is to simply take derivatives in the matrix problem. Indeed, the objective Eq. (1) is convex in X, and the problem is unconstrained. Taking the matrix derivative of $\|AX - I_m\|_F^2$, and setting it equal to zero, we have

$$2A^T AX - 2A^T I_m = 0.$$

When A is full column rank, $A^T A$ is full rank and invertible, so we have $X^* = (A^T A)^{-1} A^T$.

(c) Show that in general, a solution is $X^* = A^\dagger$, the pseudo-inverse of A. *Hint:* use the SVD of A, and exploit the fact that the Frobenius norm of a matrix is the Euclidean norm of the vector formed with its singular values.

Solution: We know that if x_i^* is optimal for the ordinary least-squares problem

$$\min_x \ \|Ax - e_i\|_2,$$

then $X^* := [x_1^*, \ldots, x_m^*]$ is optimal for the original problem.

Since $x_i^* = A^\dagger e_i$ is optimal for the above problem, we obtain that $X^* = [x_1^*, \ldots, x_m^*] = A^\dagger$ is optimal for the matrix problem.

(d) Assume that we would like to form an estimate X of the pseudo-inverse of a matrix A that is as sparse as possible, while maintaining a good accuracy, expressed in terms of the objective function of problem (1). How would you formulate the corresponding trade-off? How would you use CVX to plot a trade-off curve?

Solution: Ideally, to enforce sparsity, we sould like to solve a problem of the form

$$\min_X \|AX - I_m\|_F + \lambda \operatorname{card}(X) \tag{2}$$

where $\operatorname{card}(X)$ denotes the cardinality (number of non-zeros) of X. This is, however, wildly non-convex, so we relax Eq. (2) to the ℓ_1-regularized problem

$$\min_X \|AX - I_m\|_F + \lambda \sum_{i,j} |X_{ij}|. \tag{3}$$

We can then increase or decrease λ to look at the trade-off in the error $\|AX^\star - I_m\|_F$ in terms of sparsity in $X^\star$.[1] CVX code is below, and in Fig. 1 we plot the error in $\|AX - I_m\|_F$ versus the sparsity in X.

```
m = 20;
n = 10;
A = randn(m, n);
lambdas = [0 .02 .04 .08 .16 .32];
cvx_quiet(true);
Xs = cell(numel(lambdas), 1);
for i = 1:numel(lambdas);
  lambda = lambdas(i);
  cvx_begin
    variable X(n, m);
    minimize(norm(A*X - eye(m), 'fro') + lambda * sum(sum(abs(X))));
  cvx_end
  Xs{i} = X.*(abs(X) > 1e-8);
end

splevel = zeros(numel(lambdas), 1);
errlevel = zeros(numel(lambdas), 1);
for i = 1:numel(lambdas)
```

[1] Note that if A is full row-rank (so that $AX = I_m$ is solvable), it may make sense to minimize $\sum_{i,j} |X_{ij}|$ subject to $AX = I_m$.

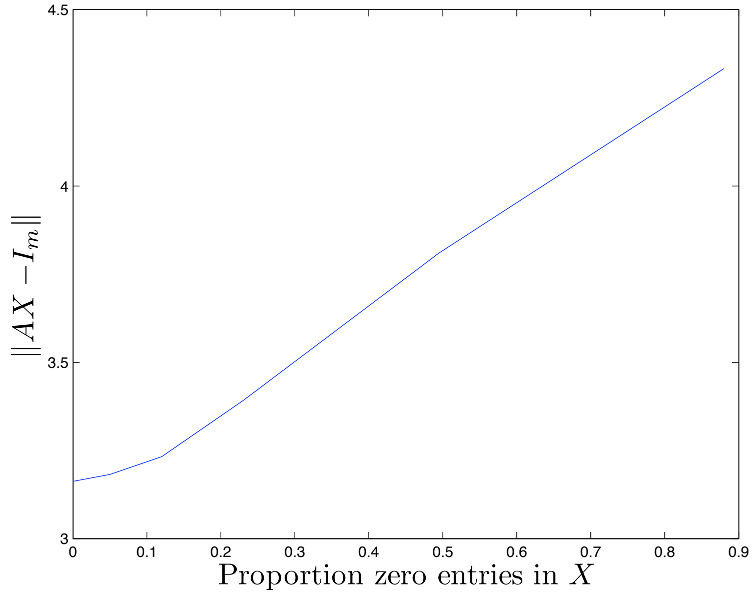

Figure 1: Tradeoffs for sparsity in terms of reconstruction error.

```
  splevel(i) = (n * m - nnz(Xs{i})) / (m * n);
  errlevel(i) = norm(A*Xs{i} - eye(m), 'fro');
end

figure;
plot(splevel, errlevel);
xlabel('Proportion zero entries in $X$', ...
       'Interpreter', 'latex', 'fontsize', 20);
h = ylabel('$\|AX - I_m\|$');
set(h, 'Interpreter', 'latex', 'fontsize', 20);
```

3. (*3 points*) Let $p_0, p_1, \ldots, p_m$ be a collection of $(m+1)$ points in $\mathbf{R}^n$. Consider the set of points closer (in Euclidean norm) to p_0 than to the remaining points $p_1, \ldots, p_m$:

$$\mathcal{P} = \{x \in \mathbf{R}^n \ : \ \forall\, i = 1, \ldots, m, \ \ \|x - p_0\|_2 \le \|x - p_i\|_2\}\,.$$

Show that the set $\mathcal{P}$ is a polyhedron, and provide a representation of it in terms of the problem's data, as $\mathcal{P} = \{x \ : \ Ax \le b\}$, where A is matrix and b a vector, which you will determine.

Solution: There are (at least) two equally reasonable ways to solve this problem. The first is to simply re-write the constraints and re-arrange symbols to give ourselves polyhedral inequalities. That is, we note that

$$\left\|x - p_0\right\|_2 \le \left\|x - p_i\right\|_2 \quad \text{iff} \quad \left\|x - p_0\right\|_2^2 \le \left\|x - p_i\right\|_2^2,$$

and the second inequality can be re-written as

$$x^T x - 2p_0^T x + p_0^T p_0 \le x^T x - 2p_i^T x + p_i^T p_i \quad \Leftrightarrow \quad 2p_i^T x - 2p_0^T x \le p_i^T p_i - p_0^T p_0$$
$$\Leftrightarrow \quad (p_i - p_0)^T x \le \frac{1}{2}(p_i - p_0)^T (p_i + p_0)$$

Evidently, then, if we set $A \in \mathbf{R}^{m \times n}$ and b as

$$A = \begin{bmatrix} (p_1 - p_0)^T \\ (p_2 - p_0)^T \\ \vdots \\ (p_m - p_0)^T \end{bmatrix} \quad \text{and} \quad b = \frac{1}{2}\begin{bmatrix} (p_1 - p_0)^T (p_1 + p_0) \\ (p_2 - p_0)^T (p_2 + p_0) \\ \vdots \\ (p_m - p_0)^T (p_m + p_0) \end{bmatrix}, \tag{4}$$

we can write the set as $\mathcal{P} = \{x : Ax \le b\}$.

A more geometric way of deriving the above is as follows. First, we recognize that we will be intersecting a series of convex sets, so if we can describe $\|x - p_0\|_2 \le \|x - p_i\|_2$ as a linear inequality, we can simply stack the inequalities. Now, the set of points that is closer to the point p_0 than it is to the point p_i is a half-space, and the normal to the half-space is $p_i - p_0$. Now all that remains is to find the point on the line halfway between p_0 and p_i, because this will be the point that gives the inequality defining our half-space. Clearly, this point is $(p_i + p_0)/2$. See Fig. 2 for a picture. Thus, since our normal is given by $(p_i - p_0)$ and we know that all points in the half space must be "closer" to p_0 than the point $(p_i + p_0)/2$, we have all x in the half space satisfy $(p_i - p_0)^T x \le (p_i - p_0)^T (p_i + p_0)/2$, which gives the same linear inequalities as the matrix A and vector b in Eq. (4).

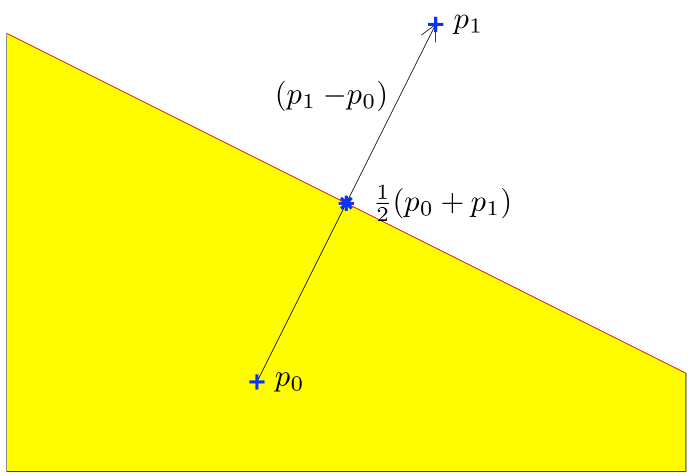

Figure 2: Halfspace separation

4. (*3 points*) We consider a production planning problem. The variables $x_j \in \mathbf{R}$, $j = 1, \ldots, n$, in the problem are activity levels (for example, production levels for different products manufactured by a company). These activities consume m resources, which

are limited. Activity j consumes $A_{ij}x_j$ of resource i. (Ordinarily we have $A_{ij} \geq 0$, that is, activity j *consumes* resource i, but we allow the possibility that $A_{ij} < 0$, which means that activity j actually generates resource i as a by-product.) The total resource consumption is additive, so the total of resource i consumed is $c_i = \sum_{j=1} A_{ij}x_j$. Each resource consumption is limited:for every activity i, we must have $c_i \leq c_i^{\max}$ where $c_i^{\max}$ are given. Activity j generates revenue $r_j(x_j)$, where r_j is a function of the form

$$r_j(u) = \begin{cases} p_j u & \text{if } 0 \leq u \leq q_j, \\ p_j q_j + p_j^{\text{disc}}(u - q_j) & \text{otherwise,} \end{cases}$$

where $q_j > 0$, $p_j^{\text{disc}} > 0$ are given parameters. The above revenue model says that above a certain level of activity, the revenue is discounted, and not as high as for lower activity level.

Show how to formulate the problem of maximizing revenue under the resource constraints as an LP.

Solution: We let our optimization variables be r and x, where r_j corresponds to the r_j above. We want to maximize r_j (our revenue) subject to the constraint that $r_j \leq \min\{p_j x_j, p_j q_j + p_j^{\text{disc}}(x_j - q_j)\}$, which is to say, r_j is less than both. If we define the matrices and vectors

$$P = \mathbf{diag}(p_1, \ldots, p_n),\ P^{\text{disc}} = \mathbf{diag}(p_1^{\text{disc}}, \ldots, p_n^{\text{disc}}),\ q = \begin{bmatrix} q_1 \\ \vdots \\ q_n \end{bmatrix},\ c^{\max} = \begin{bmatrix} c_1^{\max} \\ \vdots \\ c_n^{\max} \end{bmatrix}$$

we can re-write the problem as

$$\begin{array}{ll} \max\limits_{r,x} & 1^T r \\ \text{subject to} & Ax \preceq c^{\max},\ x \succeq 0 \\ & r \preceq Px,\ r \preceq Pq + P^{\text{disc}}(x - q). \end{array}$$

5. (*4 points*) Consider the *unconstrained* QP

$$p^* = \min_x \ \frac{1}{2} x^T Q x - c^T x$$

where $Q = Q^T \in \mathbf{R}^{n \times n}$, $Q \succeq 0$, and $c \in \mathbf{R}^n$ are given. The goal of this problem is to determine the optimal value p^* and the optimal set $\mathcal{X}^{\text{opt}}$, in terms of c and the eigenvalues and eigenvectors of the (symmetric) matrix Q.

(a) Assume that $Q \succ 0$. Show that the optimal set is a singleton, and that p^* is finite. Determine both in terms of Q, c.

Solution: The problem is convex and unconstrained. Thus, the optimal set is that of points such that the gradient of the objective is zero. Denoting the objective functionby f, the optimality condition is

$$\nabla f(x) = Qx - c.$$

Since $Q \succ 0$, Q is invertible, and there is a unique solution to $Qx = c$. Thus $x^* = Q^{-1}c$ is the unique optimal point, and

$$p^* = \frac{1}{2}(x^*)^T Q x^* - c^T x^* = \frac{1}{2} c^T Q^{-1} c - c^T Q^{-1} c = -\frac{1}{2} c^T Q^{-1} c.$$

Assume from now on that Q is not invertible.

(b) Assume further that Q is diagonal: $Q = \mathbf{diag}(\lambda_1, \ldots, \lambda_n)$, with $\lambda_1 \geq \ldots \geq \lambda_r > \lambda_{r+1} = \ldots = \lambda_n = 0$, where r is the rank of Q ($1 \leq r < n$). Solve the problem in that case. (You will distinguish between two cases.)

Solution: Partition the variable x in two parts: $x = (y, z)$ with $y \in \mathbf{R}^r$ and $z \in \mathbf{R}^{n-r}$. Likewise partition c in two parts: $c = (a, b)$ with $a \in \mathbf{R}^r$ and $b \in \mathbf{R}^{n-r}$. Since $Q = \mathbf{diag}(\Lambda, 0)$, with $\Lambda = \mathbf{diag}(\lambda_1, \ldots, \lambda_r)$, we have

$$f(x) = \frac{1}{2} x^T Q x - c^T x = \frac{1}{2} y^T \Lambda y - a^T y - b^T z = f_1(y) + f_2(z),$$

where $f_1(y) = \frac{1}{2} y^T \Lambda y - a^T y$ and $f_2(z) = b^T z$.

The above is a sum of two functions f_1, f_2 of independent variables y and z. Thus our problem can be solved by minimizing f_1 over y and f_2 over z independently:

$$\min_{x=(y,z)} \; f_1(y) + f_2(z) = \min_y \; f_1(y) + \min_z \; f_2(z).$$

The minimization over y is exactly the same as in the previous part, since the matrix Λ is positive definite. That is, the optimal y is unique, and given by $y^* = \Lambda^{-1} a$.

For the minimization with respect to z, we distinguish two cases. Either the vector $b = 0$, in which case the optimal value of $\min_z \; f_2(z)$ is zero, and *any* z is optimal. Or, $b \neq 0$, in which case the corresponding optimal value is $-\infty$ and the problem is not attained.

To summarize:

- If $b = 0$, then the optimal set is the subspace

$$\mathcal{X}^{\text{opt}} = \left\{ (\Lambda^{-1} a, z) \;:\; z \in \mathbf{R}^{n-r} \right\},$$

  and the optimal value is $-\frac{1}{2} a^T \Lambda^{-1} a$.
- If $b \neq 0$, the optimal value is $-\infty$, and the optimal set is empty (the optimal value is not attained).

Note that the condition $b = 0$ can simply be expressed as $c \in \mathbf{R}(Q)$, where $\mathbf{R}(Q)$ is the range of Q.

(c) Now we do not assume that Q is diagonal anymore. Under what conditions (on Q, c) is the optimal value finite?

Solution: Let $Q = U\tilde{\Lambda}U^T$ be the eigenvalue decomposition of the symmetric matrix Q, with U orthogonal and $\tilde{\Lambda} = \mathbf{diag}(\Lambda, 0)$, and $\Lambda \succ 0$ diagonal. Our problem reads

$$\min_x \frac{1}{2}x^T U^T \tilde{\Lambda} U x - c^T x = \min_{\tilde{x}} \frac{1}{2}\tilde{x}^T \tilde{\Lambda}\tilde{x} - \tilde{c}^T\tilde{x},$$

where $\tilde{x} := Ux$, and $\tilde{c} = Uc$. This change of variable allows us to go back to the diagonal case, which we solved in the previous part.

Partition $\tilde{c} = (\tilde{a}, \tilde{b})$ with $\tilde{a} \in \mathbf{R}^r$ the component of $\tilde{c}$ along the range of $\tilde{\Lambda}$, and $\tilde{b} \in \mathbf{R}^{n-r}$ the component of $\tilde{c}$ orthogonal to the range of $\tilde{\Lambda}$. We obtain that if $\tilde{b} \neq 0$, then the optimal value is $-\infty$ and the optimal set is empty. This corresponds to the case when c is not in the range of Q.

We conclude that the optimal value is finite if and only if c is in the range of Q.

(d) Determine the optimal value and optimal set. Be as specific as you can.

Solution: We present two solutions.

a) Assume that c is in the range of Q, so that the optimal value is finite. If that is the case, then c is also in the range of $Q^{1/2}$, because the ranges of Q and its symmetric square root are identical (both matrices share the same system of eigenvectors, and the range depends only on the eigenvectors).

Let d be such that $c = Q^{1/2}d$. We can write our objective function, f, as

$$f(x) = \frac{1}{2}x^T Q x - c^T x = \frac{1}{2}x^T Q x - d^T Q^{1/2} x = \frac{1}{2}\|Q^{1/2}x - d\|_2^2 - \|d\|_2^2.$$

Minimizing f amounts to solve the ordinary least-squares problem

$$\min_x \|Q^{1/2}x - d\|_2.$$

The solution set is

$$\mathcal{X}^{\text{opt}} = (Q^{1/2})^\dagger d + \mathbf{N}(Q^{1/2}),$$

where $\mathbf{N}(Q^{1/2})$ is the range of $Q^{1/2}$ is the range of $Q^{1/2}$, which, as noted before is the same as the range of Q itself. We can write the vector $(Q^{1/2})^\dagger d$ in terms of Q, c as follows:

$$(Q^{1/2})^\dagger d = Q^\dagger c.$$

Indeed, if $Q = U^T \tilde{\Lambda} U$ is an eigenvalue decomposition of Q, with $\tilde{\Lambda} = \mathbf{diag}(\Lambda, 0)$, and $\Lambda \succ 0$ diagonal, then

$$(Q^{1/2})^\dagger d = U^T \begin{pmatrix} \Lambda^{-1/2} & 0 \\ 0 & 0 \end{pmatrix} U d = \underbrace{U^T \begin{pmatrix} \Lambda^{-1} & 0 \\ 0 & 0 \end{pmatrix} U}_{=Q^\dagger} \underbrace{U^T \begin{pmatrix} \Lambda^{1/2} & 0 \\ 0 & 0 \end{pmatrix} d}_{=c} = Q^\dagger c.$$

Hence, if $c \in \mathbf{Range}(Q)$, we have

$$\mathcal{X}^{\text{opt}} = Q^\dagger c + \mathbf{N}(Q),$$

where $\mathbf{N}(Q)$ is the nullspace of Q. The optimal value is the value of the objective evaluated at any of the points in the optimal set. For example,

$$p^* = \frac{1}{2}(Q^\dagger c)^T Q (Q^\dagger c) - c^T Q^\dagger c = -\frac{1}{2} c^T Q^\dagger c,$$

where we have exploited the fact that $Q^\dagger Q Q^\dagger = Q^\dagger$.

b) Here is an alternative solution. We take the same notation as in the previous part. Assume that $\tilde{b} = 0$, that is, if c is in the range of Q, so that the optimal value is finite. Then the optimal set (in the space of $\tilde{x}$-variables) is

$$\tilde{\mathcal{X}} = \left\{ (\Lambda^{-1}\tilde{a}, z) \ : \ z \in \mathbf{R}^{n-r} \right\}.$$

For every $\tilde{x} \in \tilde{\mathcal{X}}$, we can write

$$\tilde{x} = \tilde{x}_0 + \begin{pmatrix} 0 \\ z \end{pmatrix},$$

where

$$\tilde{x}_0 = \begin{pmatrix} \Lambda^{-1} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \tilde{a} \\ \tilde{b} \end{pmatrix} = \begin{pmatrix} \Lambda^{-1} & 0 \\ 0 & 0 \end{pmatrix} \tilde{c}$$

With $\tilde{c} = Uc$:

$$x_0 := U^T \tilde{x}_0 = U^T \begin{pmatrix} \Lambda^{-1} & 0 \\ 0 & 0 \end{pmatrix} Uc = Q^\dagger c.$$

When z spans $\mathbf{R}^{n-r}$, the vector $(0, z)$ spans the nullspace of $\mathbf{diag}(\Lambda, 0)$, so that the vector $U^T(0, z)$ spans the nullspace of Q.

Hence, if $c \in \mathbf{Range}(Q)$, the optimal set is

$$\mathcal{X}^{\text{opt}} = Q^\dagger c + \mathbf{N}(Q),$$

where $\mathbf{N}(Q)$ is the nullspace of Q.