Math 484: Nonlinear Programming[1] Mikhail Lavrov

Chapter 4, Lecture 4: Minimum Norm Solutions

February 27, 2019 *University of Illinois at Urbana-Champaign*

1 The problem

Up until now, we have been looking at the problem of approximately solving an overconstrained system: when $A\mathbf{x} = \mathbf{b}$ has no solutions, finding an $\mathbf{x}$ that is the closest to being a solution, by minimizing $\|A\mathbf{x} - \mathbf{b}\|$.

Today, we go on to consider the opposite case: systems of equations $A\mathbf{x} = \mathbf{b}$ with infinitely many solutions. For such equations, we want to find the solution with the smallest norm, solving the optimization problem

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \quad & \|\mathbf{x}\| \\ \text{subject to} \quad & A\mathbf{x} = \mathbf{b}. \end{aligned}$$

This problem looks different, but it is also a minimum distance problem. We need to find the element of

$$S = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$$

closest to $\mathbf{0}$.

In fact, we can reduce the problem of finding a minimum-norm solution to an instance of the problem we've already solved. This will be a roundabout way to go, and in the end, we'll prove a theorem about how to get the answer directly.

2 Applying the least-squares technique

To turn the problem into one we already know how to solve, the first step is to understand the set $S = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$.

This kind of set is sometimes called an *affine subspace*: it looks just like a vector subspace of $\mathbb{R}^n$, except that it does not contain $\mathbf{0}$. From a linear algebra class, you may remember that we may write S as $S' + \mathbf{x}^{(0)}$, where:

- $\mathbf{x}^{(0)}$ is an arbitrary element of S: one particular vector in $\mathbb{R}^n$ satisfying $A\mathbf{x}^{(0)} = \mathbf{b}$.
- S' is the solution set of the corresponding homogeneous equation: $S' = \{\mathbf{y} \in \mathbb{R}^n : A\mathbf{y} = \mathbf{0}\}$.

  (In particular, S' is a vector subspace of $\mathbb{R}^n$.)
- The notation $S' + \mathbf{x}^{(0)}$ means that we add $\mathbf{x}^{(0)}$ to every element of S'.

So if we take the problem "find the element of S closest to $\mathbf{0}$" and translate it by the vector $-\mathbf{x}^{(0)}$, we turn it into a problem "find the element of S' closest to $-\mathbf{x}^{(0)}$", which is a problem we already

[1]This document comes from the Math 484 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/484-spring-2019.html

know how to solve. If we take the solution to the new problem, and translate it back by $\mathbf{x}^{(0)}$, we solve the original problem.

This strategy is illustrated in the diagram below. The original problem is in black, and the transformed problem is red. The point we're trying to get closest to is marked by a ■ and the optimal solution by a ◎.

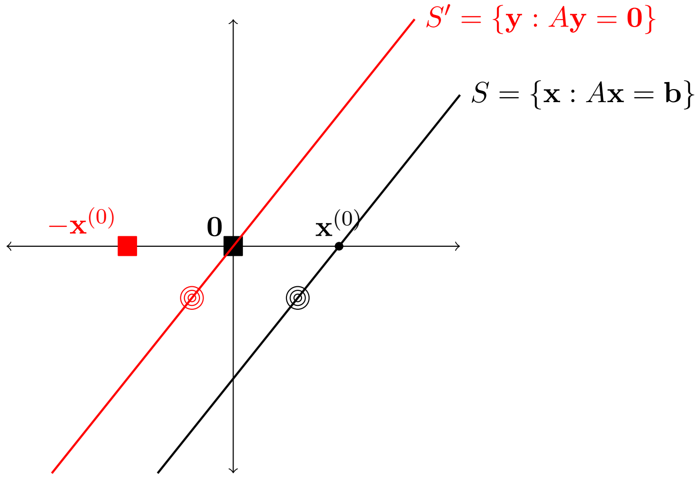

2.1 Example

What we've done doesn't give us the full understanding of the problem yet, but we can already use this idea to solve an example problem:

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^4}{\text{minimize}} \quad & \|\mathbf{x}\| \\ \text{subject to} \quad & 2x_1 - x_2 + x_3 - x_4 = 3, \\ & x_2 - x_3 - x_4 = 1. \end{aligned}$$

By the usual Gaussian elimination, we can write down a generic solution to the system of equations. If we add the second equation to the first, we get $2x_1 - 2x_4 = 4$, or $x_1 - x_4 = 2$. This lets us solve for x_1 and x_2 in terms of $x_3 - x_4$: $x_1 = 2 + x_4$, and $x_2 = 1 + x_3 + x_4$. Therefore a generic solution is

$$\mathbf{x} = \begin{bmatrix} 2 + x_4 \\ 1 + x_3 + x_4 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 0 \\ 0 \end{bmatrix} + x_3 \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} + x_4 \begin{bmatrix} 1 \\ 1 \\ 0 \\ 1 \end{bmatrix}.$$

So we've found a parametrization of the solution space in terms of a specific solution $\mathbf{x}^{(0)} = (2, 1, 0, 0)$, plus a description of the solutions to the homogeneous system $A\mathbf{y} = \mathbf{0}$.

After translating by $-\mathbf{x}^{(0)} = (-2, -1, 0, 0)$, our problem becomes to find the linear combination of $(0, 1, 1, 0)$ and $(1, 1, 0, 1)$ closest to $\mathbf{x}^{(0)}$: the least-squares minimization

$$\underset{(x_3,x_4)\in\mathbb{R}^2}{\text{minimize}} \left\| \begin{bmatrix} 0 & 1 \\ 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_3 \\ x_4 \end{bmatrix} - \begin{bmatrix} -2 \\ -1 \\ 0 \\ 0 \end{bmatrix} \right\|.$$

From what we already know, we can solve this by solving the system

$$\begin{bmatrix}0 & 1 & 1 & 0\\1 & 1 & 0 & 1\end{bmatrix}\begin{bmatrix}0 & 1\\1 & 1\\1 & 0\\0 & 1\end{bmatrix}\begin{bmatrix}x_3\\x_4\end{bmatrix} = \begin{bmatrix}0 & 1 & 1 & 0\\1 & 1 & 0 & 1\end{bmatrix}\begin{bmatrix}-2\\-1\\0\\0\end{bmatrix} \iff \begin{bmatrix}2 & 1\\1 & 3\end{bmatrix}\begin{bmatrix}x_3\\x_4\end{bmatrix} = \begin{bmatrix}-1\\-3\end{bmatrix}$$

which leads us to the solution $(x_3, x_4) = (0, -1)$. Substituting these values for x_3 and x_4 gives the solution $(x_1, x_2, x_3, x_4) = (1, 0, 0, -1)$ to the original problem.

3 The shortcut method

There is a simpler approach.

We know that if we're trying to find the closest point $\mathbf{y}^* \in S'$ to the point $-\mathbf{x}^{(0)}$, then the vector from $-\mathbf{x}^{(0)}$ to $\mathbf{y}^*$ will be perpendicular to S'. Going back to the diagram:

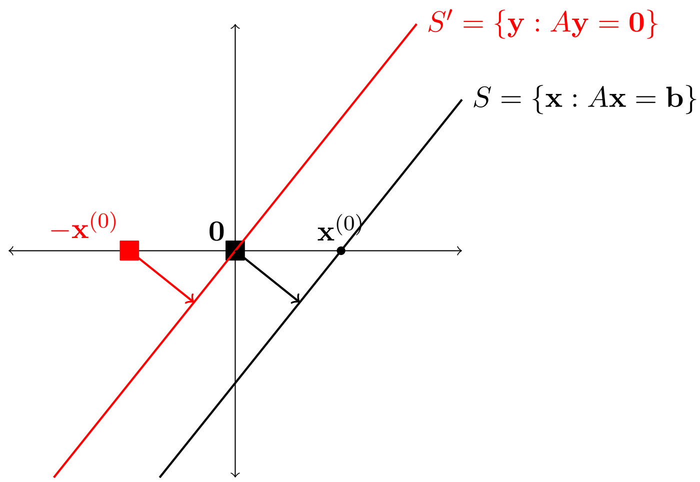

We know from last week that the red arrow is perpendicular to S'. This is preserved by translation, so the *black* arrow—which represents the vector $\mathbf{x}^*$ we're trying to find—is also perpendicular to S'.

So we have a perpendicularity condition in the minimum-norm problem as well:

Lemma 3.1. *A vector $\mathbf{x}^*$ satisfying $A\mathbf{x}^* = \mathbf{b}$ is the minimum-norm solution to the system of equations $A\mathbf{x} = \mathbf{b}$ if and only if $\mathbf{x}^* \cdot \mathbf{y} = 0$ for all solutions $\mathbf{y}$ of the homogeneous system $A\mathbf{y} = \mathbf{0}$.*

There's another way to phrase this condition. From linear algebra, we know that the null space and the row space of a matrix are orthogonal complements. So $\mathbf{x}^*$ is orthogonal to the null space of A (to all vectors $\mathbf{y}$ such that $A\mathbf{y} = \mathbf{0}$) precisely when $\mathbf{x}^*$ is an element of the row space of A: ($\mathbf{x}^* = A^\mathsf{T}\mathbf{w}$ for some $\mathbf{w}$). For a proof of this, in case it was not covered in a linear algebra class you took, see the next page.

Using this fact, we have:

Theorem 3.1. *A vector $\mathbf{x}^*$ satisfying $A\mathbf{x}^* = \mathbf{b}$ is the minimum-norm solution to the system of equations $A\mathbf{x} = \mathbf{b}$ if and only if it can be written as $\mathbf{x}^* = A^\mathsf{T}\mathbf{w}$ for some $\mathbf{w}$.*

So the minimum-norm solution $\mathbf{x}^*$ *can be found by solving the system*

$$AA^\mathsf{T}\mathbf{w} = \mathbf{b}$$

for $\mathbf{w}$*, then setting* $\mathbf{x}^* = A^\mathsf{T}\mathbf{w}$*.*

Using this, finding the solution is much faster. Going back to the example problem

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^4}{\text{minimize}} \quad & \|\mathbf{x}\| \\ \text{subject to} \quad & 2x_1 - x_2 + x_3 - x_4 = 3, \\ & x_2 - x_3 - x_4 = 1. \end{aligned}$$

we write down the system

$$\begin{bmatrix} 2 & -1 & 1 & -1 \\ 0 & 1 & -1 & -1 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ -1 & 1 \\ 1 & -1 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix} \iff \begin{bmatrix} 7 & -1 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$$

which gives us $(w_1, w_2) = (\frac{1}{2}, \frac{1}{2})$ and the same final solution $(x_1, x_2, x_3, x_4) = (1, 0, 0, -1)$.

The orthogonal complement of the null space

For the sake of completeness—in case you did not see it in a linear algebra class, or have forgotten—here is a proof that the row space and the null space of A are orthogonal complements. I've stated this result here in exactly the form we need to use it in the lecture.

Theorem 3.2. *Given an $m \times n$ matrix A, the following two conditions are equivalent for a vector $\mathbf{x}^* \in \mathbb{R}^n$:*

(a) We can write $\mathbf{x}^*$ as $A^\mathsf{T}\mathbf{w}$ for some $\mathbf{w} \in \mathbb{R}^m$.

(b) For all $\mathbf{y} \in \mathbb{R}^n$ such that $A\mathbf{y} = \mathbf{0}$, we have $\mathbf{x}^* \cdot \mathbf{y} = \mathbf{0}$.

Proof. First, let's suppose condition (a) holds for $\mathbf{x}^*$: we *can* write $\mathbf{x}^*$ as $A^\mathsf{T}\mathbf{w}$. This means that for all $\mathbf{y}$ such that $A\mathbf{y} = \mathbf{0}$, we have

$$\mathbf{x}^* \cdot \mathbf{y} = \mathbf{x}^{*\mathsf{T}}\mathbf{y} = (A^\mathsf{T}\mathbf{w})^\mathsf{T}\mathbf{y} = \mathbf{w}^\mathsf{T}A\mathbf{y} = \mathbf{w}^\mathsf{T}\mathbf{0} = 0$$

so condition (b) also holds for $\mathbf{x}^*$.

Second, let's suppose condition (a) does not hold $\mathbf{x}^*$: we *cannot* write $\mathbf{x}^*$ as $A^\mathsf{T}\mathbf{w}$. What does this mean? It means that, taking $A^\mathsf{T}\mathbf{w} = \mathbf{x}^*$ as an equation where $\mathbf{w}$ is the unknown, the equation has no solution.

This can only happen because we can deduce a contradiction from $A^\mathsf{T}\mathbf{w} = \mathbf{x}^*$: by adding up the rows of A^T with some coefficients $y_1, y_2, \dots, y_n$, we get a row of zeroes, and by adding up the entries of $\mathbf{x}^*$ with the same coefficients, we get a nonzero value.

Writing this down in matrix language, we get $\mathbf{y}^\mathsf{T}A^\mathsf{T} = \mathbf{0}^\mathsf{T}$ and $\mathbf{y}^\mathsf{T}\mathbf{x}^* \neq 0$. In other words, there is a vector $\mathbf{y} \in \mathbb{R}^n$ such that $A\mathbf{y} = \mathbf{0}$, but $\mathbf{x}^* \cdot \mathbf{y} \neq 0$. So condition (b) also does not hold for $\mathbf{x}^*$. □