5 Least Squares Problems

Consider the solution of $A\boldsymbol{x} = \boldsymbol{b}$, where $A \in \mathbb{C}^{m \times n}$ with $m > n$. In general, this system is *overdetermined* and no exact solution is possible.

Example Fit a straight line to 10 measurements. If we represent the line by $f(x) = mx + c$ and the 10 pieces of data are $\{(x_1, y_1), \ldots, (x_{10}, y_{10})\}$, then the constraints can be sumamrized in the linear system

$$\underbrace{\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_{10} & 1 \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} m \\ c \end{bmatrix}}_{\boldsymbol{x}} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{10} \end{bmatrix}}_{\boldsymbol{b}}.$$

This type of problem is known as *linear regression* or (linear) *least squares fitting.*

The basic idea (due to Gauss) is to minimize the 2-norm of the *residual vector*, i.e.,

$$\|\boldsymbol{b} - A\boldsymbol{x}\|_2.$$

In other words, we want to find $\boldsymbol{x} \in \mathbb{C}^n$ such that

$$\sum_{i=1}^{m} \left[\boldsymbol{b}_i - (A\boldsymbol{x})_i\right]^2$$

is minimized. For the notation used in the example above we want to find m and c such that $\sum_{i=1}^{10} \left[y_i - (mx_i + c)\right]^2$ is minimized.

Example We can generalize the previous example to polynomial least squares fitting of arbitrary degree. To this end we assume that

$$p(x) = \sum_{i=0}^{n} c_i x^i,$$

where n is the degree of the polynomial.

We can fit a polynomial of degree n to $m > n$ data points (x_i, y_i), $i = 1, \ldots, m$, using the least squares approach, i.e.,

$$\min \sum_{i=1}^{m} \left[y_i - p(x_i)\right]^2$$

is used as constraint for the overdetermined linear system $A\boldsymbol{x} = \boldsymbol{b}$ with

$$A = \begin{bmatrix} 1 & x_1 & x_1^2 & \ldots & x_1^n \\ 1 & x_2 & x_2^2 & \ldots & x_2^n \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_m & x_m^2 & \ldots & x_m^n \end{bmatrix}, \quad \boldsymbol{x} = \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix}.$$

Remark The special case $n = m - 1$ is called *interpolation* and is known to have a unique solutions if the conditions are independent, i.e., the points x_i are distinct. However, for large degrees n we frequently observe severe oscillations which is undesirable.

5.1 How to Compute the Least Squares Solution

We want to find $\boldsymbol{x}$ such that $A\boldsymbol{x} \in \text{range}(A)$ is as close as possible to a given vector $\boldsymbol{b}$.

It should be clear that we need $A\boldsymbol{x}$ to be the orthogonal projection of $\boldsymbol{b}$ onto the range of A, i.e.,

$$A\boldsymbol{x} = P\boldsymbol{b}.$$

Then the residual $\boldsymbol{r} = \boldsymbol{b} - A\boldsymbol{x}$ will be minimal.

Theorem 5.1 *Let $A \in \mathbb{C}^{m\times n}$ $(m \geq n)$ with* rank$(A) = n$ *and $\boldsymbol{b} \in \mathbb{C}^m$. The vector $\boldsymbol{x} \in \mathbb{C}^n$ minimizes the residual norm $\|\boldsymbol{r}\|_2 = \|\boldsymbol{b} - A\boldsymbol{x}\|_2$*

$$\begin{aligned} \textit{if and only if} \qquad A^*\boldsymbol{r} &= \boldsymbol{0}, && (20) \\ \textit{if and only if} \qquad A^*A\boldsymbol{x} &= A^*\boldsymbol{b}, && (21) \\ \textit{if and only if} \qquad A\boldsymbol{x} &= P\boldsymbol{b}, && (22) \end{aligned}$$

where $P \in \mathbb{C}^{m\times m}$ is the orthogonal projector onto the range of A. Moreover, A^*A is nonsingular and the least squares solution $\boldsymbol{x}$ is unique.

(20) says that $\boldsymbol{r}$ is perpendicular to the range of A. (21) is known as the set of *normal equations.*

Proof To see that (20) $\Leftrightarrow$ (21) we use the definition of the residual $\boldsymbol{r} = \boldsymbol{b} - A\boldsymbol{x}$. Then

$$\begin{aligned} A^*\boldsymbol{r} = 0 \qquad &\Longleftrightarrow \qquad A^*(\boldsymbol{b} - A\boldsymbol{x}) = 0 \\ &\Longleftrightarrow \qquad A^*A\boldsymbol{x} = A^*\boldsymbol{b}. \end{aligned}$$

To see that (21) $\Leftrightarrow$ (22) we use that the orthogonal projector onto the range of A is given by

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

Then

$$\begin{aligned} P\boldsymbol{b} = A\boldsymbol{x} \qquad &\Longleftrightarrow \qquad A(A^*A)^{-1}A^*\boldsymbol{b} = A\boldsymbol{x} \\ &\Longleftrightarrow \qquad \underbrace{A^*A(A^*A)^{-1}}_{=I}A^*\boldsymbol{b} = A^*A\boldsymbol{x} \\ &\Longleftrightarrow \qquad A^*A\boldsymbol{x} = A^*\boldsymbol{b}. \end{aligned}$$

Note that A^*A is nonsingular if and only if A has full rank n. ■

Remark If A has full rank then A^*A is also Hermitian positive definite, i.e., $\boldsymbol{x}^*A^*A\boldsymbol{x} > 0$ for any nonzero n-vector $\boldsymbol{x}$.

For full-rank A we can take (21) and obtain the least squares solution as

$$\boldsymbol{x} = \underbrace{(A^*A)^{-1}A^*}_{=A^+}\boldsymbol{b}.$$

The matrix A^+ is known as the *pseudoinverse* of A.

5.2 Algorithms for finding the Least Squares Solution

5.2.1 Cholesky Factorization

This can be applied for a full-rank matrix A. As mentioned above A^*A is Hermitian positive definite and one can apply a symmetric form of Gaussian elimination resulting in

$$A^*A = R^*R$$

with *upper triangular* matrix R (more details will be provided in a later section).

This means that we have

$$\begin{aligned} A^*A\boldsymbol{x} = A^*\boldsymbol{b} \quad &\iff \quad R^*R\boldsymbol{x} = A^*\boldsymbol{b} \\ &\iff \quad R^*\boldsymbol{w} = A^*\boldsymbol{b} \end{aligned}$$

with $\boldsymbol{w} = R\boldsymbol{x}$. Since R is upper triangular (and R^* is lower triangular) this is easy to solve.

We obtain one of our three-step algorithms:

Algorithm (Cholesky Least Squares)

(0) Set up the problem by computing A^*A and $A^*\boldsymbol{b}$.

(1) Compute the Cholesky factorization $A^*A = R^*R$.

(2) Solve the lower triangular system $R^*\boldsymbol{w} = A^*\boldsymbol{b}$ for $\boldsymbol{w}$.

(3) Solve the upper triangular system $R\boldsymbol{x} = \boldsymbol{w}$ for $\boldsymbol{x}$.

The operations count for this algorithm turns out to be $\mathcal{O}(mn^2 + \frac{1}{3}n^3)$.

Remark The solution of the normal equations is likely to be unstable. Therefore this method is *not recommended* in general. For small problems it is usually safe to use.

5.2.2 QR Factorization

This works also for full-rank matrices A. Recall that the reduced QR factorization is given by $A = \hat{Q}\hat{R}$ with $\hat{Q}$ an $m \times n$ matrix with orthonormal columns, and $\hat{R}$ an $n \times n$ upper triangular matrix.

Now the normal equations can be re-written as

$$A^*A\boldsymbol{x} = A^*\boldsymbol{b} \quad \iff \quad \hat{R}^* \underbrace{\hat{Q}^*\hat{Q}}_{=I} \hat{R}\boldsymbol{x} = \hat{R}^*\hat{Q}^*\boldsymbol{b}.$$

Since A has full rank $\hat{R}$ will be invertible and we can further simplify to

$$\hat{R}\boldsymbol{x} = \hat{Q}^*\boldsymbol{b}.$$

This is only one triangular system to solve. The algorithm is

Algorithm (QR Least Squares)

(0) Set up the problem by computing A^*A and $A^*\boldsymbol{b}$.

(1) Compute the reduced QR factorization $A = \hat{Q}\hat{R}$.

(2) Compute $\hat{Q}^*\boldsymbol{b}$.

(3) Solve the upper triangular system $\hat{R}\boldsymbol{x} = \hat{Q}^*\boldsymbol{b}$ for $\boldsymbol{x}$.

An alternative interpretation is based on condition (22) in Theorem 5.1. We take $A = \hat{Q}\hat{R}$ and $P = \hat{Q}\hat{Q}^*$ (since $\hat{Q}$ is an orthonormal basis for range(A)). Then we have

$$\hat{Q}\hat{R}\boldsymbol{x} = \hat{Q}\hat{Q}^*\boldsymbol{b}.$$

Multiplication by $\hat{Q}^*$ yields

$$\underbrace{\hat{Q}^*\hat{Q}}_{=I}\hat{R}\boldsymbol{x} = \underbrace{\hat{Q}^*\hat{Q}}_{=I}\hat{Q}^*\boldsymbol{b}$$

so that we have

$$\hat{R}\boldsymbol{x} = \hat{Q}^*\boldsymbol{b}$$

as before.

From either interpretation we see that

$$\boldsymbol{x} = \hat{R}^{-1}\hat{Q}^*\boldsymbol{b}$$

so that (with this notation) the pseudoinverse is given by

$$A^+ = \hat{R}^{-1}\hat{Q}^*.$$

This is well-defined since $\hat{R}^{-1}$ exists because A has full rank.

The operations count (using Householder reflectors to compute the QR factorization) is $\mathcal{O}(2mn^2 - \frac{2}{3}n^3)$.

Remark This approach is more stable than the Cholesky approach and is considered the *standard method for least squares problems.*

5.2.3 SVD

We again assume that A has full rank. Recall that the reduced SVD is given by $A = \hat{U}\hat{\Sigma}V^*$, where $\hat{U} \in \mathbb{C}^{m\times n}$, $\hat{\Sigma} \in \mathbb{R}^{n\times n}$, and $V \in \mathbb{C}^{n\times n}$.

We start again with the normal equations

$$\begin{aligned} A^*A\boldsymbol{x} = A^*\boldsymbol{b} \quad &\iff \quad V\underbrace{\hat{\Sigma}^*}_{=\hat{\Sigma}}\underbrace{\hat{U}^*\hat{U}}_{=I}\hat{\Sigma}V^*\boldsymbol{x} = V\hat{\Sigma}\hat{U}^*\boldsymbol{b} \\ &\iff \quad V\hat{\Sigma}^2V^*\boldsymbol{x} = V\hat{\Sigma}\hat{U}^*\boldsymbol{b}. \end{aligned}$$

Since A has full rank we can invert $\hat{\Sigma}$ and multiply the last equation by $\hat{\Sigma}^{-1}V^*$. This results in

$$\hat{\Sigma}V^*\boldsymbol{x} = \hat{U}^*\boldsymbol{b}$$

or

$$\hat{\Sigma}\boldsymbol{w} = \hat{U}^*\boldsymbol{b} \qquad \text{with } \boldsymbol{w} = V^*\boldsymbol{x}.$$

Therefore (with the SVD notation) the pseudoinverse is given by

$$A^+ = V\hat{\Sigma}^{-1}U^*$$

and the least squares solution is given by

$$\boldsymbol{x} = A^+\boldsymbol{b} = V\hat{\Sigma}^{-1}U^*\boldsymbol{b}.$$

The algorithm is

Algorithm (SVD Least Squares)

(1) Compute the reduced SVD $A = \hat{U}\hat{\Sigma}V^*$.

(2) Compute $\hat{U}^*\boldsymbol{b}$.

(3) Solve the diagonal system $\hat{\Sigma}\boldsymbol{w} = \hat{U}^*\boldsymbol{b}$ for $\boldsymbol{w}$.

(4) Compute $\boldsymbol{x} = V\boldsymbol{w}$.

This time the operations count is $\mathcal{O}(2mn^2 + 11n^3)$ which is comparable to that of the QR factorization provided $m \gg n$. Otherwise this algorithm is more expensive, but also more stable.

5.3 Solution of Rank Deficient Least Squares Problems

If $\text{rank}(A) < n$ (which is possible even if $m < n$, i.e., if we have an *underdetermined* problem), then infinitely many solutions exist.

A common approach to obtain a well-defined solution in this case is to add an additional constraint of the form

$$\|\boldsymbol{x}\| \longrightarrow \min,$$

i.e., we seek the *minimum norm solution.*

In this case the unique solution is given by

$$\boldsymbol{x} \sum_{\sigma_i \neq 0} \frac{\boldsymbol{u}_i^*\boldsymbol{b}}{\sigma_i}\boldsymbol{v}_i$$

or

$$\boldsymbol{x} = V\Sigma^+U^*\boldsymbol{b},$$

where now the pseudoinverse is given by

$$A^+ = V\Sigma^+U^*.$$

Here the pseudoinverse of Σ is defined as

$$\Sigma^+ = \begin{bmatrix} \Sigma_1^{-1} & O \\ O & O \end{bmatrix},$$

where Σ_1 is that part of Σ containing the positive (and therefore invertible) singular values.

As a final remark we note that there exists also a variant of the QR factorization that is more stable due to the use of column pivoting. The idea of pivoting will be discussed later in the context of the LU factorization.