6 Least Squares Approximation by QR Factorization

6.1 Formulation of Least Squares Approximation Problems

Least-squares problems arise, for instance, when one seeks to determine the relation between an independent variable, say time, and a measured dependent variable, such as position or velocity of an object.

Example 6.1

When we drop a ball a few feet above the ground with initial speed zero, it will fall towards the ground. In the absence of air resistance, its speed is known to be a linear function of time (because the acceleration is constant). However, measurements of the speed of the ball at different times during its fall to the ground might not be on a straight line, due to air resistance, sudden wind bursts, and inaccuracies in the measurements. Typically, we are interested in determining the relation between time and speed in the absence of measurement errors and wind bursts. □

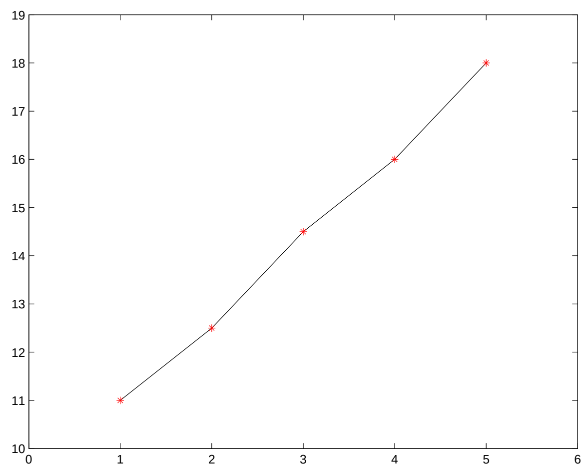

Figure 1: Example 6.2: The data points $\{x_j, y_j\}$, $1 \leq j \leq 5$, are marked by red crosses. The graph shown is a piecewise linear function that connects the data points.

Example 6.2

The pressure in a fully automatic espresso machine varies with the size of the coffee grinds. The pressure is measured in bar and the grind size in units, where 1 signifies a coarse grind and 5 a fine one. The relation is shown in Figure 1. The graph suggests that the relation between the grind size x and the pressure y might be linear with the measured pressure (the y_j) contaminated by measurement error. □

The present lecture is concerned with the determination of relations of the form

$$y = f(x),$$

[0]Version October 2, 2013

where x is the independent variable, y is the dependent variable, and f is a function, which defines this dependence. Let $x_1, x_2, \ldots, x_m$ be known values of the dependent variable, and let $y_1^{\text{exact}}, y_2^{\text{exact}}, \ldots, y_m^{\text{exact}}$ be the associated, but *unknown*, values of the function f, i.e.,

$$y_j^{\text{exact}} = f(x_j), \qquad 1 \leq j \leq m.$$

We are interested in the common situation when only approximations, y_j, of the values y_j^{exact} are available. Thus,

$$y_j = y_j^{\text{exact}} + \eta_j, \qquad 1 \leq j \leq m,$$

where the η_j are errors, which may stem from careless measurement, flaws in the measuring devise, or inaccuracies introduced during transmission of the data from the measurement devise to the computer.

In many applications, the unavailable error-free data y_j^{exact}, $1 \leq j \leq m$, can be well approximated by the values of a linear combination of a few simple functions, such as monomials, exponential functions, or trigonometric functions at the x_j. For instance, Figure 1 suggests the relation between the x_j and y_j^{exact} to be linear. In order to gain insight into the relation depicted by the figure, we seek to fit the linear polynomial

$$p(x) = c_1 + c_2 x \tag{1}$$

to the available data $\{x_j, y_j\}_{j=1}^5$ of Example 6.2. Knowing the coefficients c_1 and c_2 sheds light on the relation between the grind size and pressure. We refer to the linear polynomial as our *model*, because we use it to model the relation between the grind size and pressure.

There are many ways to determine a linear polynomial that fits the data in Example 6.2 in some sense. We will determine the polynomial that minimizes the least-squares error,

$$\min_{p \text{ linear}} \sum_{j=1}^{m} (p(x_j) - y_j)^2. \tag{2}$$

Minimizing the least-squares error can be justified statistically when the data errors η_j are independent and normally distributed with zero mean. These conditions often are met, at least roughly, in applications. Moreover, there are fairly fast numerical methods available for the solution of least-squares minimization problems. We will discuss several of these methods in this course.

Substituting the expression (1) into (2) transforms the least-squares minimization problem into the form

$$\min_{c_1, c_2} \left\| \begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_{m-1} \\ 1 & x_m \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} - \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{m-1} \\ y_m \end{bmatrix} \right\|^2. \tag{3}$$

Thus, we seek to determine the linear combination of the columns of the matrix that best approximates the data-vector in the least-squares sense. We note that the square in (3) can be removed without changing the solution of the least-squares problem.

Least-squares minimization problems arise in many applications and do not have to be related to the approximation of a function by monomials or other functions. Let the vector $\mathbf{y} = [y_1, y_2, \ldots, y_m]^T \in \mathbb{R}^m$ contain available, possibly error-contaminated, data, and assume that we would like to fit $n < m$ linearly independent vectors $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n \in \mathbb{R}^m$ to $\mathbf{y}$. We may formulate this task as a least-squares problem as follows. Introduce the matrix

$$A = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n] \in \mathbb{R}^{m \times n} \tag{4}$$

and solve the least-squares minimization problem

$$\min_{\mathbf{c} \in \mathbb{R}^n} \|A\mathbf{c} - \mathbf{y}\|, \qquad \mathbf{c} = [c_1, c_2, \ldots, c_n]^T. \tag{5}$$

Throughout this section, we will assume that the columns of A are linearly independent. Then $\text{null}(A) = \{0\}$ and the least-squares problem has a unique solution. The following sections describe numerical methods for the solution of least-squares problems.

6.2 Solution of Least-Squares Problems by QR Factorization

When the matrix A in (5) is upper triangular with zero padding, the least-squares problem can be solved by back substitution. This is illustrated in the following example.

Example 6.3

Let

$$\hat{R} = \begin{bmatrix} R \\ O \end{bmatrix} \in \mathbb{R}^{m \times n}, \qquad m > n, \tag{6}$$

where $R \in \mathbb{R}^{n \times n}$ is a nonsingular upper triangular matrix and $O \in \mathbb{R}^{(m-n) \times n}$ is a matrix with all entries zero. Let $\mathbf{y} \in \mathbb{R}^m$. The solution of the least-squares problem

$$\min_{\mathbf{c} \in \mathbb{R}^n} \|\hat{R}\mathbf{c} - \mathbf{y}\| \tag{7}$$

is given by $\mathbf{c} = R^{-1}\mathbf{y}_{1:n}$ and can be computed by back substitution. Here $\mathbf{y}_{1:n}$ denotes the vector made up of the first n entries of $\mathbf{y}$. Thus,

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_{1:n} \\ \mathbf{y}_{(n+1):m} \end{bmatrix}.$$

The vector $\mathbf{y}_{(n+1):m}$ does not affect the solution $\mathbf{c}$ of the least-squares problem. Nevertheless, the norm of the vector $\mathbf{y}_{(n+1):m}$ is of interest. This will be commented on in Example 6.4. □

We will describe how to factor a general $m \times n$ matrix A, with $m \geq n$, into an orthogonal matrix $\hat{Q} \in \mathbb{R}^{m \times m}$ and a matrix of the form (6), i.e.,

$$A = \hat{Q}\hat{R}. \tag{8}$$

Since the columns of A are linearly independent, so are the columns of $\hat{R}$. In particular, the square submatrix R is nonsingular.

The factorization (8) differs from the QR factorization (4) of Lecture 5 in that the matrix $\hat{Q}$ above is square also when A is not. Both the factorizations (4) of Lecture 5 and (8) above are referred to as *QR factorizations* of A. They both can be used to solve least-square problems of the form (5). This is discussed below. The factorization (8) above also is essential in the QR algorithm for computing eigenvalues and eigenvectors of a matrix. The QR algorithm will be discussed in a later lecture. Here we only note that it requires QR factorization (8) of a sequence of square matrices.

Before discussing the computation of the QR factorization (8), we comment on its usefulness for the solution of least-squares problems. Substitute the QR factorization (8) into the least-squares problem (5) to obtain

$$\min_{\mathbf{c}\in\mathbb{R}^n} \|\hat{Q}\hat{R}\mathbf{c}-\mathbf{y}\|.$$

Since the norm of the vector $\hat{Q}\hat{R}\mathbf{c}-\mathbf{y}$ does not change by multiplication by an orthogonal matrix, see (5) of Lecture 5, it follows that

$$\|Q\hat{R}\mathbf{c}-\mathbf{y}\| = \|\hat{Q}^T(\hat{Q}\hat{R}\mathbf{c}-\mathbf{y})\| = \|\hat{R}\mathbf{c}-\hat{Q}^T\mathbf{y}\|$$

and, therefore,

$$\min_{\mathbf{c}\in\mathbb{R}^n} \|A\mathbf{c}-\mathbf{y}\| = \min_{\mathbf{c}\in\mathbb{R}^n} \|\hat{Q}\hat{R}\mathbf{c}-\mathbf{y}\| = \min_{\mathbf{c}\in\mathbb{R}^n} \|\hat{R}\mathbf{c}-\hat{Q}^T\mathbf{y}\|. \tag{9}$$

Thus, we compute the solution of the minimization problem on the left-hand side by solving the minimization problem on the right-hand side. The latter is of the form (7) and can be solved by back substitution.

Example 6.4

The matrix $\hat{R}$ in (9) is of the form (6). Similarly as in Example 6.3, we let R denote the leading $n\times n$ submatrix of $\hat{R}$. Define the vector $\hat{\mathbf{y}} = \hat{Q}^T\mathbf{y}$ in (9). The solution of (9) is given by $\mathbf{c} = R^{-1}\mathbf{y}_{1:n}$; cf. Example 6.3. Here we comment on the role of the trailing entries of the vector $\mathbf{y}$.

For definiteness, let A be the $m\times 2$ matrix in the least-squares problem (3) and consider (9). Then $n=2$ and the solution $\mathbf{c}\in\mathbb{R}^2$ determines the linear combination of the columns of A that best approximates the data $\mathbf{y}\in\mathbb{R}^m$. The entries in $\hat{\mathbf{y}}_{3:m}$ is the part of the data that cannot be approximated by the linear model (2). If the norm $\|\hat{\mathbf{y}}_{3:m}\|$ is small, then the model can represent the data well. There may be two reasons for $\|\hat{\mathbf{y}}_{3:m}\|$ to be large: there are large errors in the data $\mathbf{y}$ but the chosen model is appropriate, or the model is poorly suited to approximate the data. In the latter case, another model should possibly be used. For instance, one might consider fitting a quadratic polynomial instead of a linear one. We refer to the norm $\|\mathbf{y}_{(n+1):m}\|$ as the *discrepancy.* □

The next section describes how to compute the factorization (8) with the aid of Householder matrices and illustrates the application of this factorization to the solution of least-squares problems. A later section discusses how the QR factorization (5) of Lecture 5 can be used to solve least-squares problem. A difficulty that has to be addressed is that the matrix Q in the latter factorization is not square and, therefore, not an orthogonal matrix.

6.3 QR Factorization by Householder Matrices

We describe the computation of the QR factorization (8) of $A\in\mathbb{R}^{m\times n}$, $m\geq n$, with the aid of Householder matrices,

$$H = I-\rho\mathbf{u}\mathbf{u}^T \in \mathbb{R}^{m\times m}, \qquad \mathbf{u}\neq\mathbf{0}, \qquad \rho = \frac{2}{\mathbf{u}^T\mathbf{u}}, \tag{10}$$

introduced in Lecture 5. For notational simplicity, we let $n=3$ in our discussion. Let $\mathbf{a}_j$ denote the jth columns of the matrix A, i.e.,

$$A = [\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3] \in \mathbb{R}^{m\times 3}.$$

Let $H^{(1)} = H$ be the Householder matrix that maps the column $\mathbf{a}_1$ onto $\sigma^{(1)}\mathbf{e}_1$, where $\sigma^{(1)}$ is a scalar and, as usual, $\mathbf{e}_1$ denotes the first axis vector. Then

$$H^{(1)}A = [H^{(1)}\mathbf{a}_1, H^{(1)}\mathbf{a}_2, H^{(1)}\mathbf{a}_3] = [\sigma^{(1)}\mathbf{e}_1, H^{(1)}\mathbf{a}_2, H^{(1)}\mathbf{a}_3].$$

The first column of the above matrix is the first column of the upper triangular matrix $\hat{R}$ in the QR factorization (8). We turn to the computation of the second column of this matrix.

Let the matrix $A^{(2)} = [\mathbf{a}_1^{(2)}, \mathbf{a}_2^{(2)}] \in \mathbb{R}^{(m-1)\times 2}$ be made up of the entries in rows 2 through m and columns 2 and 3 of $H^{(1)}A$. Thus, we obtain $A^{(2)}$ from $H^{(1)}A$ by removing in the first row and the first column. Let $H^{(2)} \in \mathbb{R}^{(m-1)\times(m-1)}$ denote the Householder matrix that maps the vector $\mathbf{a}_1^{(2)}$ onto a multiple of $\mathbf{e}_1$, i.e.,

$$H^{(2)}\mathbf{a}_1^{(2)} = \sigma^{(2)}\mathbf{e}_1,$$

where $\sigma^{(2)}$ is a suitable scalar. Multiplying $A^{(2)}$ by $H^{(2)}$ from the left-hand side yields

$$H^{(2)}A^{(2)} = [H^{(2)}\mathbf{a}_1^{(2)}, H^{(2)}\mathbf{a}_2^{(2)}] = [\sigma^{(2)}\mathbf{e}_1, H^{(2)}\mathbf{a}_2^{(2)}].$$

Prepend a new first row and a new first column to $H^{(2)}$ to obtain the matrix

$$\hat{H}^{(2)} = \begin{bmatrix} 1 & \mathbf{0}^T \\ \mathbf{0} & H^{(2)} \end{bmatrix} \in \mathbb{R}^{m\times m} \tag{11}$$

of the same size as $H^{(1)}$. Here $\mathbf{0} \in \mathbb{R}^{(m-1)}$ denotes a (column) vector with all entries zero. It is easy to verify that $\hat{H}^{(2)}$ is orthogonal; see Exercise 6.2.

Multiplication of $H^{(1)}A$ by $\hat{H}^{(2)}$ leaves the first row and the first column of $H^{(1)}A$ unchanged. (Verify this!) Therefore,

$$\hat{H}^{(2)}H^{(1)}A = \begin{bmatrix} \sigma^{(1)} & * & * \\ 0 & \sigma^{(2)} & * \\ & 0 & * \\ & & \\ \vdots & \vdots & \vdots \\ & & \\ 0 & 0 & * \end{bmatrix} \in \mathbb{R}^{m\times 3}. \tag{12}$$

The entries 2 through m of the last columns make up the vector $H^{(2)}\mathbf{a}_2^{(2)}$. Entries marked by $*$ may be nonvanishing.

One more Householder transformation has to be applied in order to bring the matrix (12) into upper triangular form. Let the vector $\mathbf{a}_1^{(3)} \in \mathbb{R}^{(m-2)}$ be made up of the entries 3 through m of the last column of the matrix (12), and let $H^{(3)} \in \mathbb{R}^{(m-2)\times(m-2)}$ denote the Householder matrix that maps $\mathbf{a}_1^{(3)}$ onto a multiple of $\mathbf{e}_1$, i.e.,

$$H^{(3)}\mathbf{a}_1^{(3)} = \sigma^{(3)}\mathbf{e}_1.$$

Prepend 2 new first rows and 2 new first columns to $H^{(3)}$ to obtain the matrix

$$\hat{H}^{(3)} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & & & \\ \vdots & \vdots & & H^{(3)} & \\ 0 & 0 & & & \end{bmatrix} \in \mathbb{R}^{m\times m} \tag{13}$$

of the same size as $H^{(1)}$. We may think of $\hat{H}^{(3)}$ as being an $m \times m$ identity matrix, with the trailing principal $(m-2)\times(m-2)$ submatrix replaced by $H^{(3)}$. The matrix $\hat{H}^{(3)}$ is orthogonal.

Analogously to (12), we obtain

$$\hat{R} = \hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}A = \begin{bmatrix} \sigma^{(1)} & * & * \\ 0 & \sigma^{(2)} & * \\ & 0 & \sigma^{(3)} \\ & & 0 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 0 \end{bmatrix} \in \mathbb{R}^{m\times 3}. \tag{14}$$

Multiplying (14) from the left-hand side by $\hat{H}^{(3)}$, $\hat{H}^{(2)}$, and $H^{(1)}$, in order, and using that

$$\hat{H}^{(j)}\hat{H}^{(j)} = (\hat{H}^{(j)})^T\hat{H}^{(j)} = I, \qquad j = 2, 3,$$

as well as a similar formula for $H^{(1)}$, we obtain

$$A = H^{(1)}\hat{H}^{(2)}\hat{H}^{(3)}\hat{R}. \tag{15}$$

Since the product of orthogonal matrices is orthogonal, see Exercise 5.6 of Lecture 5, the matrix $\hat{Q} = H^{(1)}\hat{H}^{(2)}\hat{H}^{(3)}$ is orthogonal and (15) is a QR factorization of the form (8).

In order to solve the least-squares minimization problem on the right-hand side of (9), we also have to compute $\hat{Q}^T\mathbf{y}$. Note that

$$\hat{Q}^T\mathbf{y} = \hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}\mathbf{y},$$

i.e., we can apply the matrices $H^{(1)}$, $\hat{H}^{(2)}$, and $\hat{H}^{(3)}$ to $\mathbf{y}$ in the order they are generated to obtain the vector $\hat{Q}^T\mathbf{y}$. In particular, the entries of the matrix $\hat{Q}$ do not have to be computed explicitly.

An algorithm is said to be *backward stable* if it in the presence of round-off errors introduced during the computations determines the exact solution to a nearby problem. QR factorization with the aid of Householder matrices is backward stable. Let $\text{fl}(\hat{Q})$ and $\text{fl}(\hat{R})$ denote the computed factors in finite precision arithmetic, where we represent $\text{fl}(\hat{Q})$ as a product of unmultiplied Householder-type matrices. Each Householder matrix is represented by one vector. Thus, neither $\text{fl}(\hat{Q})$ nor the Householder matrices are explicitly formed. Then one can show that

$$\frac{\|\text{fl}(\hat{Q})\text{fl}(\hat{R}) - A\|}{\|A\|} = \mathcal{O}(\mathsf{eps}), \tag{16}$$

where $\mathcal{O}(t)$ denotes a quantity bounded by $c|t|$ for some constant $c \geq 0$ independent of t, as t approaches zero. Note that $\text{fl}(\hat{Q})$ is numerically orthogonal, since perturbations in the vectors that represent the Householder matrices do not change the orthogonality property. Moreover, $\text{fl}(\hat{R})$ is upper triangular by construction.

Example 6.5

Table 1 shows measured function values y at four times t. Figure 2 displays the data of the table. The figure suggests that, in the absence of measurement errors, the data points might lie on a parabola, i.e., y may be a quadratic function of t. We therefore would like to determine the quadratic polynomial

$$p(t) = c_1 + c_2t + c_3t^2 \tag{17}$$

t	y
1	1.0
2	1.5
3	3.0
4	6.0

Table 1: Example 6.5: The right-hand side column displays measured values at several times t.

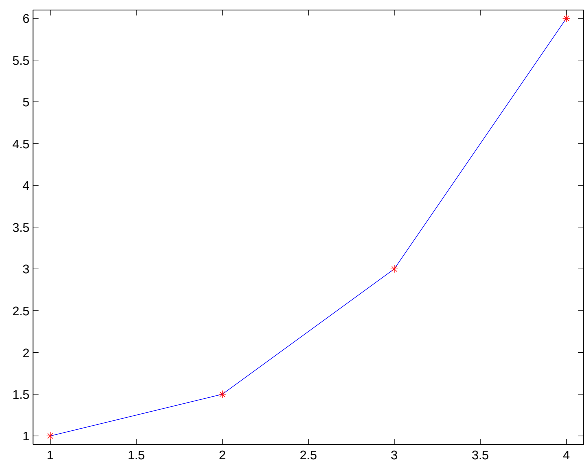

Figure 2: Example 6.5: Graph for Table 2. The data points $\{x_j, y_j\}$, $1 \leq j \leq 4$, are marked by red crosses. The graph shown is a piecewise linear function, which connects the data points. The graph indicates that the relation plotted may be quadratic.

that best approximates the data of Table 1 in the least-squares sense. Tabulation of the polynomial at the nodes $t_j = j$, $1 \leq j \leq 4$, gives the matrix $A \in \mathbb{R}^{4\times 3}$. We denote the vector with the available function values by $\mathbf{y}$. Thus, we have

$$A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \\ 1 & 4 & 16 \end{bmatrix}, \qquad \mathbf{y} = \begin{bmatrix} 1.0 \\ 1.5 \\ 3.0 \\ 6.0 \end{bmatrix}.$$

We would like to solve the minimization problem

$$\min_{\mathbf{c}\in\mathbb{R}^3} \|A\mathbf{c} - \mathbf{y}\|, \qquad \mathbf{c} = [c_1, c_2, c_3]^T. \tag{18}$$

This least-squares problem is solved by application of a sequence of judiciously chosen Householder matrices, which we apply from the left-hand side to both A and $\mathbf{y}$. The Householder matrices are designed to transform A into upper triangular form. The first Householder matrix $H^{(1)}$ is chosen to map the first column of A

onto a multiple of $\mathbf{e}_1$, i.e., we would like

$$H^{(1)} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \sigma^{(1)} \mathbf{e}_1$$

for some scalar $\sigma^{(1)}$. This Householder matrix is the same one as in Example 5.6 of Lecture 5. Thus, $\sigma^{(1)} = -2$ and we obtain, after rounding to 2 decimals,

$$H^{(1)} A = \begin{bmatrix} -2 & -5 & -15 \\ 0 & 0 & -\ 1.33 \\ 0 & 1 & 3.67 \\ 0 & 2 & 10.67 \end{bmatrix}, \qquad H^{(1)} \mathbf{y} = \begin{bmatrix} -5.75 \\ -0.75 \\ 0.75 \\ 3.75 \end{bmatrix}.$$

Our next task is to determine a Householder matrix

$$H^{(2)} = I - \rho^{(2)} \mathbf{u}^{(2)} (\mathbf{u}^{(2)})^T \in \mathbb{R}^{3 \times 3} \tag{19}$$

that maps the vector consisting of the entries 2 through 4 of column 2 of the matrix $H^{(1)} A$ onto a multiple of $\mathbf{e}_1$, i.e., we would like $H^{(2)}$ to be such that

$$H^{(2)} \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix} = \sigma^{(2)} \mathbf{e}_1$$

for some scalar $\sigma^{(2)}$. Then $|\sigma^{(2)}| = \|[0, 1, 2]^T\| = \sqrt{5} \approx 2.24$. Since the first component of the vector to be mapped vanishes, the sign of $\sigma^{(2)}$ can be chosen arbitrarily; we let $\sigma^{(2)} = \sqrt{5}$ and obtain

$$u^{(2)} = [-\sqrt{5}, 1, 2]^T.$$

This yields $\rho^{(2)} = 2/\|\mathbf{u}^{(2)}\|^2 = 1/5$ and it follows from (19) that

$$H^{(2)} = \begin{bmatrix} 0 & 0.45 & 0.89 \\ 0.45 & 0.80 & -0.40 \\ 0.89 & -0.40 & 0.20 \end{bmatrix}.$$

We embed this matrix in the 4×4 orthogonal matrix

$$\hat{H}^{(2)} = \begin{bmatrix} 1 & \mathbf{0}^T \\ \mathbf{0} & H^{(2)} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0.45 & 0.89 \\ 0 & 0.45 & 0.80 & -0.40 \\ 0 & 0.89 & -0.40 & 0.20 \end{bmatrix},$$

where $\mathbf{0}$ denotes the zero vector in $\mathbb{R}^3$. Then

$$\hat{H}^{(2)} H^{(1)} A = \begin{bmatrix} -2 & -5.00 & -15.00 \\ 0 & 2.24 & 11.18 \\ 0 & 0 & -1.93 \\ 0 & 0 & -0.53 \end{bmatrix}, \qquad \hat{H}^{(2)} H^{(1)} \mathbf{y} = \hat{H}^{(2)} \begin{bmatrix} -5.75 \\ -0.75 \\ 0.75 \\ 3.75 \end{bmatrix} = \begin{bmatrix} -5.75 \\ 3.69 \\ -1.24 \\ -0.22 \end{bmatrix}.$$

Thus, we have determined the first 2 columns of the upper triangular matrix $\hat{R}$ in (8). We remark that it is not necessary to compute the entries of $\hat{H}^{(2)}$ in order to evaluate the above matrix and vector.

The next Householder matrix

$$H^{(3)} = I - \rho^{(3)}\mathbf{u}^{(3)}(\mathbf{u}^{(3)})^T \in \mathbb{R}^{2\times 2}$$

is designed to zero the last entry in the last column of the matrix $\hat{H}^{(2)}H^{(1)}A$; it is determined by the last two entries of this column. Thus, we would like $H^{(3)}$ to satisfy

$$H^{(3)}\begin{bmatrix} -1.93 \\ -0.53 \end{bmatrix} = \sigma^{(3)}\mathbf{e}_1$$

for some scalar $\sigma^{(3)}$. We obtain $|\sigma^{(3)}| = \|[-1.93, -0.53]^T\| = 2$ and choose $\sigma^{(3)} = 2$ in order to avoid loss of accuracy due to cancellation of significant digits. Thus,

$$\mathbf{u}^{(3)} = \begin{bmatrix} -1.93 - \sigma^{(3)} \\ -0.53 \end{bmatrix} = \begin{bmatrix} -3.93 \\ -0.53 \end{bmatrix}.$$

It follows that $\rho^{(3)} = 0.13$ and

$$H^{(3)} = \begin{bmatrix} -0.96 & -0.26 \\ -0.26 & 0.96 \end{bmatrix}.$$

We embed $H^{(3)}$ in the larger orthogonal matrix

$$\hat{H}^{(3)} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -0.96 & -0.26 \\ 0 & 0 & -0.26 & 0.96 \end{bmatrix}$$

and finally obtain

$$\hat{R} = \hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}A = \begin{bmatrix} -2 & -5.00 & -15.00 \\ 0 & 2.24 & 11.18 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}, \qquad \hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}\mathbf{y} = \hat{H}^{(3)}\begin{bmatrix} -5.75 \\ 3.69 \\ -1.24 \\ -0.22 \end{bmatrix} = \begin{bmatrix} -5.75 \\ 3.69 \\ 1.25 \\ 0.11 \end{bmatrix}.$$

Let R be the leading 3×3 submatrix of $\hat{R}$ and let the vector $\mathbf{d}$ consist of the first 3 entries of the vector above, i.e.,

$$R = \begin{bmatrix} -2 & -5.00 & -15.00 \\ 0 & 2.24 & 11.18 \\ 0 & 0 & 2 \end{bmatrix}, \qquad \mathbf{d} = \begin{bmatrix} -5.75 \\ 3.69 \\ 1.25 \end{bmatrix}.$$

The solution $\mathbf{c} = [c_1, c_2, c_3]^T = [1.875, -1.475, 0.625]^T$ of the linear system of equations $R\mathbf{c} = \mathbf{d}$ solves the least-squares problem (18).

Figure 3 shows the graph of Figure 2 and the graph of the quadratic polynomial (17) determined by the computed coefficients c_j. The polynomial can be seen to approximate the data well, which indicates that modeling by a quadratic polynomial may be appropriate.

The last entry, 0.11, of the vector $\hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}\mathbf{y}$ is the discrepancy. It provides a numerical value of the goodness of fit, because

$$\min_{\mathbf{c}\in\mathbb{R}^3} \|A\mathbf{c} - \mathbf{y}\| = \min_{\mathbf{c}\in\mathbb{R}^3} \|\hat{R}\mathbf{c} - \hat{H}^{(3)}\hat{H}^{(2)}H^{(1)}\mathbf{y}\| = 0.11.$$

Since this value is small, the computed polynomial (17) provides a good fit of the data. □

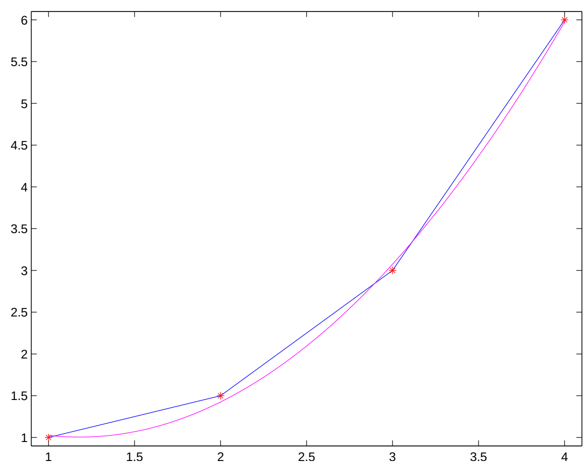

Figure 3: Example 6.5: Graph of Figure 2 and of the computed quadratic polynomial (17) (in magenta). The polynomial is seen to fit the data well.

6.4 QR Factorization by the Modified Gram-Schmidt Procedure

Above we used Householder matrices to compute the QR factorization (8) and the applied this factorization to the solution of least-squares problems. Here we will show that how the QR factorization (4) of Lecture 5, determined by modified Gram-Schmidt procedure, can be applied to the solution of least-squares problems (5). In exact arithmetic, also a QR factorization determined by the standard Gram-Schmidt procedure can be used; however, this is not advisable when the factorization is determined using floating point arithmetic.

Split the matrix $\hat{Q}$ in the QR factorization (8) into two parts:

$$\hat{Q} = [Q, \breve{Q}], \qquad Q \in \mathbb{R}^{m\times n}, \qquad \breve{Q} \in \mathbb{R}^{m\times(m-n)}. \tag{20}$$

Then, using (6), we obtain

$$A = \hat{Q}\hat{R} = [Q, \breve{Q}] \begin{bmatrix} R \\ O \end{bmatrix} = QR.$$

The right-hand side is a factorization of the form (4) of Lecture 5. This shows that orthogonalization by the modified Gram-Schmidt procedure gives the first n columns, up to factors ± 1, of the matrix $\hat{Q}$ determined by Householder matrices; see Exercise 5.2 of Lecture 5. Substituting (6) and (20) into (9) yields

$$\min_{\mathbf{c}\in\mathbb{R}^n} \|A\mathbf{c} - \mathbf{y}\|^2 = \min_{\mathbf{c}\in\mathbb{R}^n} \|\hat{R}\mathbf{c} - \hat{Q}^T\mathbf{y}\|^2 = \min_{\mathbf{c}\in\mathbb{R}^n} \{\|R\mathbf{c} - Q^T\mathbf{y}\|^2 + \|\breve{Q}^T\mathbf{y}\|^2\}.$$

The last term is independent of $\mathbf{c}$. Therefore, the solution of the above least-squares problem can be determined by solving

$$R\mathbf{c} = Q^T\mathbf{y}$$

by back substitution. Note that the matrix R above is square.

In summary, we may solve least-squares problems (5) by using the QR factorization of A determined by orthogonalization of the the columns of the matrix by the modified Gram-Schmidt procedure. This approach does not give the discrepancy $\|\breve{Q}^T\mathbf{y}\|$ during the computations. The discrepancy provides insight into how well the model fits the data.

6.5 The Normal Equations

A natural, but inherently flawed approach, to solving least-squares problems is to form and solve the so-called normal equations. Since this approach is quite popular, we provide an outline and comment on its deficiency. Substituting the expression (1) into (2) yields

$$\min_{c_1,c_2} \sum_{j=1}^{m} (c_1 + c_2 x_j - y_j)^2 . \tag{21}$$

Define the function

$$F(c_1, c_2) = \sum_{j=1}^{m} (c_1 + c_2 x_j - y_j)^2 . \tag{22}$$

The minimization problem (21) can be expressed as

$$\min_{c_1,c_2} F(c_1, c_2).$$

Keeping c_2 fixed, we see that $c_1 \to F(c_1, c_2)$ is a quadratic polynomial with positive leading coefficient. (The leading coefficient is m.) The quadratic polynomial therefore has a unique minimum, which is achieved for the value of c_1 for which the derivative of $c_1 \to F(c_1, c_2)$ vanishes. This value of c_1 depends on our choice of c_2.

The derivative of $c_1 \to F(c_1, c_2)$ is commonly referred as the *partial derivative* of F with respect to c_1 and denoted by $\partial F/\partial c_1$. Thus, for any fixed coefficient c_2, we require c_1 to satisfy

$$\frac{\partial F(c_1, c_2)}{\partial c_1} = 0. \tag{23}$$

Similarly, keeping c_1 fixed, the function $c_2 \to F(c_1, c_2)$ is a quadratic polynomial, which we seek to minimize. Analogously to the discussion above, we are lead to determining a zero of the partial derivative of F with respect to c_2, i.e., we require c_2 to satisfy

$$\frac{\partial F(c_1, c_2)}{\partial c_2} = 0. \tag{24}$$

Using the definition (22) of the function F, we obtain from (23) and (24) the equations

$$\begin{aligned} \frac{\partial F(c_1, c_2)}{\partial c_1} &= 2\sum_{j=1}^{m} (c_1 + c_2 x_j - y_j) = 0, \\ \frac{\partial F(c_1, c_2)}{\partial c_2} &= 2\sum_{j=1}^{m} (c_1 + c_2 x_j - y_j)\, x_j = 0, \end{aligned}$$

which can be expressed as a linear system of equations

$$\begin{bmatrix} 2\sum_{j=1}^{m} 1 & 2\sum_{j=1}^{m} x_j \\ 2\sum_{j=1}^{m} x_j & 2\sum_{j=1}^{m} x_j^2 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} 2\sum_{j=1}^{m} y_j \\ 2\sum_{j=1}^{m} y_j x_j \end{bmatrix}. \tag{25}$$

This system is known as the *normal equations* associated with the least-squares minimization problem (2).

Let A be the matrix in (3) and let the vector $\mathbf{y}$ have the entries y_j, i.e.,

$$A = \begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_{m-1} \\ 1 & x_m \end{bmatrix}, \qquad \mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{m-1} \\ y_m \end{bmatrix}.$$

Then the least-squares problem (3) can be expressed in the form

$$\min_{\mathbf{c}\in\mathbb{R}^2} \|A\mathbf{c} - \mathbf{y}\|_2^2, \qquad \mathbf{c} = [c_1, c_2]^T,$$

and the normal equations (25) can be written as

$$A^T A\mathbf{c} = A^T\mathbf{y}. \tag{26}$$

The normal equations (26) can be solved by a variant of Gaussian elimination. However, generally the solution of the normal equations is more sensitive to errors in the data and to round-off errors introduced during the computations than the solution of the least-squares problem (5) determined by QR factorization of the matrix using Householder matrices.

Example 6.6

Consider the matrix

$$A = \begin{bmatrix} 1 & 1 \\ \delta & 0 \\ 0 & \delta \end{bmatrix}, \qquad \delta = 1 \cdot 10^{-8}$$

and let $\mathbf{y} = [1, 2, 3]^T$. Then the solution of the least squares problem

$$\min_{\mathbf{c}\in\mathbb{R}^2} \|A\mathbf{c} - \mathbf{y}\|$$

can be computed in MATLAB, while the solution of the associated normal equations (26) cannot. The latter depends on that floating point representation of $A^T A$, given by

$$\mathrm{fl}(A^T A) = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix},$$

is a singular matrix. □

The solution of the normal equations generally should be avoided also for other matrices than the one in the above example. This depends on that the condition number of the matrix $A^T A$ is the square of the condition number of A. The condition number measures the sensitivity of a problem to perturbations. It is defined in Lecture 10. A large condition number indicates that the problem solved may be sensitive to errors in the data as well as to round-off errors introduced during the computations. For many least-squares problems, the condition number of $A^T A$ is much larger than the condition number of A. The solution of these problems should not be computed by solving the normal equations.

Even though the normal equations should not be used for computation, they do provide some geometrical insight. Express (26) in the form

$$A^T(\mathbf{y} - A\mathbf{c}) = \mathbf{0}.$$

The above equation shows that the residual vector $\mathbf{r} = \mathbf{y} - A\mathbf{c}$ associated with the least-squares solution is orthogonal to the columns of A.

Exercise 6.1

Discuss the nonuniquness of the factorization (8). Hint: Consider the first n and last $m - n$ columns of Q separately. □

Exercise 6.2

Show that the matrix (11) is orthogonal. □

1	11.0
2	12.5
3	14.5
4	16.0
5	18.0

Table 2: Exercise 6.3: Grind size (left-hand side column) versus pressure (right-hand side column).

Exercise 6.3

Table 2 shows the data for Figure 1. Use this data to determine the linear polynomial (1). Use Householder matrices to compute the solution (similarly as in Example 6.5). Show the Householder matrices used and the matrices obtained after applying 1 and 2 Householder matrices to the matrix in the least-squares problem (3). Also show the data-vectors obtained by orthogonal transformation. Compute the discrepancy and plot the computed polynomial together with the data. Does the polynomial fit the data well? □

Exercise 6.4

The MATLAB command $\mathbf{c} = A\backslash\mathbf{y}$ solves the least-squares minimization problem (5). The computations are carried out similarly as in Example 6.5, i.e., they are based on Householder matrices. Determine the polynomials of degrees 0, 1, 2, and 3 that best fit the data of Exercise 6.3 in the least-squares sense. What are the discrepancies for the different degrees? □

Exercise 6.5

The measured values of a harmonic oscillator $y(t) = a\sin(\phi + \phi_0)$ are given in Table 3. We would like to determine the amplitude a and phase ϕ_0. (a) Why can least-squares minimization not be applied in a straightforward way to determine these quantities? (b) Express $y(t)$ in the form

$$y(t) = c_1\cos(\phi) + c_2\sin(\phi)$$

ϕ	y
4	3.41
34	7.70
64	9.84
94	9.40

Table 3: Exercise 6.5: The right-hand side column displays the angles in degrees and the left hand-side column the associated function values.

and determine the coefficients c_1 and c_2 by the least-squares minimization. Then seek to determine ϕ_0 from c_1 and c_2 by using a trigonometric identity. Plot the computed function $y(t)$ and data of Table 3. $\Box$

Exercise 6.6

Consider the computation of the QR factorization of $A \in \mathbb{R}^{m \times n}$. The number of arithmetic floating point operations (flops) required for computing the QR factorization is a function of the form $f(m,n)$. This function can be expanded into products of powers of m and n. What are the powers of the leading terms? That is, determine the powers $j, k \geq 0$ of the leading term $m^j n^k$. $\Box$

Exercise 6.7

Let $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ be linearly independent vectors in $\mathbb{R}^m$, with $m \geq 3$. Determine an orthonormal basis of $\text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ with the aid of Householder matrices. $\Box$

Exercise 6.8

Solve the problem of Example 6.5 by using the modified Gram-Schmidt procedure to compute the QR factorization (4) of Lecture 5, and then apply this factorization. $\Box$