

LEAST SQUARES PROBLEMS

A typical scenario in many sciences is acquisition of m numbers to describe some object that is understood to actually require only $n \ll m$ parameters. For example, m voltage measurements u_i of an alternating current could readily be reduced to three parameters, the amplitude, phase and frequency $u(t) = a \sin(\omega t + \varphi)$. Very often a simple first-degree polynomial approximation $y = ax + b$ is sought for a large data set $D = \{(x_i, y_i), i = 1, \dots, m\}$. All of these are instances of data compression, a problem that can be solved in a linear algebra framework.

1. Projection

Consider a partition of a vector space U into orthogonal subspaces $U = V \oplus W$, $V = W^\perp$, $W = V^\perp$. Within the typical scenario described above $U = \mathbb{R}^m$, $V \subset \mathbb{R}^m$, $W \subset \mathbb{R}^m$, $\dim V = n$, $\dim W = m - n$. If $\mathbf{V} = [\mathbf{v}_1 \dots \mathbf{v}_n] \in \mathbb{R}^{m \times n}$ is a basis for V and $\mathbf{W} = [\mathbf{w}_1 \dots \mathbf{w}_{m-n}] \in \mathbb{R}^{m \times (m-n)}$ is a basis for W , then $\mathbf{U} = [\mathbf{v}_1 \dots \mathbf{v}_n \mathbf{w}_1 \dots \mathbf{w}_{m-n}]$ is a basis for U . Even though the matrices \mathbf{V}, \mathbf{W} are not necessarily square, they are said to be orthonormal, in the sense that all columns are of unit norm and orthogonal to one another. Computation of the matrix product $\mathbf{V}^T \mathbf{V}$ leads to the formation of the identity matrix within \mathbb{R}^n

$$\mathbf{V}^T \mathbf{V} = \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_n^T \end{bmatrix} [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n] = \begin{bmatrix} \mathbf{v}_1^T \mathbf{v}_1 & \mathbf{v}_1^T \mathbf{v}_2 & \dots & \mathbf{v}_1^T \mathbf{v}_n \\ \mathbf{v}_2^T \mathbf{v}_1 & \mathbf{v}_2^T \mathbf{v}_2 & \dots & \mathbf{v}_2^T \mathbf{v}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{v}_n^T \mathbf{v}_1 & \mathbf{v}_n^T \mathbf{v}_2 & \dots & \mathbf{v}_n^T \mathbf{v}_n \end{bmatrix} = \mathbf{I}_n.$$

Similarly, $\mathbf{W}^T \mathbf{W} = \mathbf{I}_{m-n}$. Whereas for the square orthogonal matrix \mathbf{U} multiplication both on the left and the right by its transpose leads to the formation of the identity matrix

$$\mathbf{U}^T \mathbf{U} = \mathbf{U} \mathbf{U}^T = \mathbf{I}_m,$$

the same operations applied to rectangular orthonormal matrices lead to different results

$$\mathbf{V}^T \mathbf{V} = \mathbf{I}_n, \mathbf{V} \mathbf{V}^T = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n] \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_n^T \end{bmatrix} = \sum_{i=1}^n \mathbf{v}_i \mathbf{v}_i^T, \text{rank}(\mathbf{v}_i \mathbf{v}_i^T) = 1$$

A simple example is provided by taking $\mathbf{V} = \mathbf{I}_{m,n}$, the first n columns of the identity matrix in which case

$$\mathbf{V} \mathbf{V}^T = \sum_{i=1}^n \mathbf{e}_i \mathbf{e}_i^T = \begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{m \times m}.$$

- Applying $\mathbf{P} = \mathbf{V} \mathbf{V}^T$ to some vector $\mathbf{b} \in \mathbb{R}^m$ leads to a vector $\mathbf{r} = \mathbf{P} \mathbf{b}$ whose first n components are those of \mathbf{b} , and the remaining $m - n$ are zero. The subtraction $\mathbf{b} - \mathbf{r}$ leads to a new vector $\mathbf{s} = (\mathbf{I} - \mathbf{P}) \mathbf{b}$ that has the first components equal to zero, and the remaining $m - n$ the same as those of \mathbf{b} . Such operations are referred to as *projections*, and for $\mathbf{V} = \mathbf{I}_{m,n}$ correspond to projection onto the $\text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$.




Figure 1. Projection in \mathbb{R}^2 . The vectors $\mathbf{r}, \mathbf{s} \in \mathbb{R}^2$ have two components, but could be expressed through scaling of $\mathbf{e}_1, \mathbf{e}_2$.

Returning to the general case, the orthogonal matrices $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{V} \in \mathbb{R}^{m \times n}$, $\mathbf{W} \in \mathbb{R}^{m \times (m-n)}$ are associated with linear mappings $\mathbf{b} = \mathbf{f}(\mathbf{x}) = \mathbf{U}\mathbf{x}$, $\mathbf{r} = \mathbf{g}(\mathbf{b}) = \mathbf{P}\mathbf{b}$, $\mathbf{s} = \mathbf{h}(\mathbf{b}) = (\mathbf{I} - \mathbf{P})\mathbf{b}$. The mapping \mathbf{f} gives the components in the \mathbf{I} basis of a vector whose components in the \mathbf{U} basis are \mathbf{x} . The mappings \mathbf{g}, \mathbf{h} project a vector onto $\text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_{m-n}\}$, respectively. When \mathbf{V}, \mathbf{W} are orthogonal matrices the projections are also orthogonal $\mathbf{r} \perp \mathbf{s}$. Projection can also be carried out onto nonorthogonal spanning sets, but the process is fraught with possible error, especially when the angle between basis vectors is small, and will be avoided henceforth.

Notice that projection of a vector already in the spanning set simply returns the same vector, which leads to a general definition.

DEFINITION. The mapping is called a *projection* if $\mathbf{f} \circ \mathbf{f} = \mathbf{f}$, or if for any $\mathbf{u} \in U$, $\mathbf{f}(\mathbf{f}(\mathbf{u})) = \mathbf{f}(\mathbf{u})$. With \mathbf{P} the matrix associated \mathbf{f} , a projection matrix satisfies $\mathbf{P}^2 = \mathbf{P}$.


$$\mathbf{P} = \mathbf{V}\mathbf{V}^T$$

$$\mathbf{P}^2 = \mathbf{P}\mathbf{P} = \mathbf{V}\mathbf{V}^T\mathbf{V}\mathbf{V}^T = \mathbf{V}(\mathbf{V}^T\mathbf{V})\mathbf{V}^T = \mathbf{V}\mathbf{I}\mathbf{V}^T = \mathbf{V}\mathbf{V}^T = \mathbf{P}$$

2. Gram-Schmidt

Orthonormal vector sets $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$ are of the greatest practical utility, leading to the question of whether some such a set can be obtained from an arbitrary set of vectors $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$. This is possible for independent vectors, through what is known as the Gram-Schmidt algorithm

1. Start with an arbitrary direction \mathbf{a}_1
2. Divide by its norm to obtain a unit-norm vector $\mathbf{q}_1 = \mathbf{a}_1 / \|\mathbf{a}_1\|$
3. Choose another direction \mathbf{a}_2
4. Subtract off its component along previous direction(s) $\mathbf{a}_2 - (\mathbf{q}_1^T \mathbf{a}_2) \mathbf{q}_1$
5. Divide by norm $\mathbf{q}_2 = (\mathbf{a}_2 - (\mathbf{q}_1^T \mathbf{a}_2) \mathbf{q}_1) / \|\mathbf{a}_2 - (\mathbf{q}_1^T \mathbf{a}_2) \mathbf{q}_1\|$
6. Repeat the above



$$P_1 a_2 = (q_1 q_1^T) a_2 = q_1 (q_1^T a_2) = (q_1^T a_2) q_1$$

The above geometrical description can be expressed in terms of matrix operations as

$$A = (a_1 \ a_2 \ \dots \ a_n) = (q_1 \ q_2 \ \dots \ q_n) \begin{pmatrix} r_{11} & r_{12} & r_{13} & \dots & r_{1n} \\ 0 & r_{22} & r_{23} & \dots & r_{2n} \\ 0 & 0 & r_{33} & \dots & r_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \dots & r_{nn} \end{pmatrix} = QR,$$

equivalent to the system

$$\begin{cases} a_1 = r_{11}q_1 \\ a_2 = r_{12}q_1 + r_{22}q_2 \\ \vdots \\ a_n = r_{1n}q_1 + r_{2n}q_2 + \dots + r_{nn}q_n \end{cases}$$

The system is easily solved by *forward substitution* resulting in what is known as the (modified) *Gram-Schmidt algorithm*, transcribed below both in pseudo-code and in Julia.

Algorithm (Gram-Schmidt)

Given n vectors a_1, \dots, a_n
 Initialize $q_1 = a_1, \dots, q_n = a_n, R = I_n$
 for $i = 1$ to n
 $r_{ii} = (q_i^T q_i)^{1/2}$
 if $r_{ii} < \epsilon$ break;
 $q_i = q_i / r_{ii}$
 for $j = i+1$ to n
 $r_{ij} = q_i^T a_j; q_j = q_j - r_{ij}q_i$
 end
 end
 return Q, R

```

∴ function mgs(A)
    m, n = size(A); Q = copy(A); R = zeros(n, n)
    for i = 1:n
        R[i, i] = sqrt(Q[:, i]' * Q[:, i])
        if (R[i, i] < eps())
            break
        end
        Q[:, i] = Q[:, i] / R[i, i]
        for j = i+1:n
            R[i, j] = Q[:, i]' * A[:, j]
            Q[:, j] = Q[:, j] - R[i, j] * Q[:, i]
        end
    end
    return Q, R
end;

```

- Note that the normalization condition $\|q_{ii}\| = 1$ is satisfied by two values $\pm r_{ii}$, so results from the above implementation might give orthogonal vectors q_1, \dots, q_n of different orientations than those returned by the Octave `qr` function. The implementation provided by computational packages such as Octave contain many refinements of the basic algorithm and it's usually preferable to use these in application

By analogy to arithmetic and polynomial algebra, the Gram-Schmidt algorithm furnishes a *factorization*

$$QR = A$$

with $Q \in \mathbb{R}^{m \times n}$ with orthonormal columns and $R \in \mathbb{R}^{n \times n}$ an upper triangular matrix, known as the QR -factorization. Since the column vectors within Q were obtained through linear combinations of the column vectors of A we have

$$C(A) = C(Q) \neq C(R)$$

3. QR solution of linear algebra problems

The QR -factorization can be used to solve basic problems within linear algebra.

3.1. Transformation of coordinates

Recall that when given a vector $b \in \mathbb{R}^m$, an implicit basis is assumed, the canonical basis given by the column vectors of the identity matrix $I \in \mathbb{R}^{m \times m}$. The coordinates x in another basis $A \in \mathbb{R}^{m \times m}$ can be found by solving the equation

$$Ib = b = Ax,$$

by an intermediate change of coordinates to the orthogonal basis Q . Since the basis Q is orthogonal the relation $Q^T Q = I$ holds, and changes of coordinates from I to Q , $Qc = b$, are easily computed $c = Q^T b$. Since matrix multiplication is associative

$$b = Ax = (QR)x = Q(Rx),$$

the relations $Rx = Q^T b = c$ are obtained, stating that x also contains the coordinates of c in the basis R . The three steps are:

1. Compute the QR -factorization, $QR = A$;
2. Find the coordinates of b in the orthogonal basis Q , $c = Q^T b$;
3. Find the coordinates of x in basis R , $Rx = c$.

Since R is upper-triangular,

$$\begin{pmatrix} r_{11} & r_{12} & r_{13} & \dots & r_{1m} \\ 0 & r_{22} & r_{23} & \dots & r_{2m} \\ 0 & 0 & r_{33} & \dots & r_{3m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \dots & r_{mm} \end{pmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_{m-1} \\ x_m \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_{m-1} \\ c_m \end{bmatrix}$$

the coordinates of c in the R basis are easily found by *back substitution*.

Algorithm (Back substitution)

```

Given R upper-triangular, vectors c
for i = m down to 1
    if rii < ε break;
    xi = ci / rii
    for j = i-1 down to 1
        cj = cj - rjixi
    end
end
end
return x

```

```

∴ function bcks(R,c)
    m,n=size(R); x=zeros(m,1)
    for i=m:-1:1
        x[i]=c[i]/R[i,i]
        for j=i-1:-1:1
            c[j]=c[j]-R[j,i]*x[i]
        end
    end
    return x
end;
∴

```

- The above operations are carried out in the background by the backslash operation $A \setminus b$ to solve $A * x = b$, inspired by the scalar mnemonic $ax = b \Rightarrow x = (1/a) b$.

3.2. General orthogonal bases

The above approach for the real vector space \mathcal{R}_m can be used to determine orthogonal bases for any other vector space by appropriate modification of the scalar product. For example, within the space of smooth functions $\mathcal{C}^\infty[-1, 1]$ that can be differentiated an arbitrary number of times, the Taylor series

$$f(x) = f(0) \cdot 1 + f'(0) \cdot x + \frac{1}{2} f''(0) \cdot x^2 + \dots + \frac{1}{n!} f^{(n)}(0) \cdot x^n + \dots +$$

is seen to be a linear combination of the monomial basis $\mathbf{M} = [1 \ x \ x^2 \ \dots]$ with scaling coefficients $\{f(0), f'(0), \frac{1}{2} f''(0), \dots\}$. The scalar product

$$(f, g) = \int_{-1}^1 f(x) g(x) dx$$

can be seen as the extension to the $[-1, 1]$ continuum of a the vector dot product. Orthogonalization of the monomial basis with the above scalar product leads to the definition of another family of polynomials, known as the Legendre polynomials

$$Q_0(x) = \left(\frac{1}{2}\right)^{1/2} \cdot 1, Q_1(x) = \left(\frac{3}{2}\right)^{1/2} \cdot x, Q_2(x) = \left(\frac{5}{8}\right)^{1/2} \cdot (3x^2 - 1), Q_4(x) = \left(\frac{7}{8}\right)^{1/2} \cdot (5x^3 - 3x), \dots$$

- The Legendre polynomials are usually given with a different scaling such that $P_k(1) = 1$, rather than the unit norm condition $\|Q_k\| = (Q_k, Q_k)^{1/2} = 1$. The above results can be recovered by sampling of the interval $[-1, 1]$ at points $x_i = (i - 1) h - 1, h = 2 / (m - 1), i = 1, \dots, m$, by approximation of the integral by a Riemann sum

$$\int_{-1}^1 f(x) L_j(x) dx \cong h \sum_{i=1}^m f(x_i) L_j(x_i) = h \mathbf{f}^T \mathbf{L}_j.$$




Figure 2. Comparison of monomial basis (left) to Legendre polynomial basis (right). The “resolution” of $P_3(x)$ can be interpreted as the number of crossings of the $y = 0$ ordinate axis, and is greater than that of the corresponding monomial x^3 .

3.3. Least squares

The approach to compressing data $D = \{(x_i, y_i) | i = 1, \dots, m\}$ suggested by calculus concepts is to form the sum of squared differences between $y(x_i)$ and y_i , for example for $y(x) = a_0 + a_1 x$ when carrying out linear regression,

$$S(a_0, a_1) = \sum_{i=1}^m (y(x_i) - y_i)^2 = \sum_{i=1}^m (a_0 + a_1 x_i - y_i)^2$$

and seek (a_0, a_1) that minimize $S(a_0, a_1)$. The function $S(a_0, a_1) \geq 0$ can be thought of as the height of a surface above the $a_0 a_1$ plane, and the gradient ∇S is defined as a vector in the direction of steepest slope. When at some point on the surface if the gradient is different from the zero vector $\nabla S \neq \mathbf{0}$, travel in the direction of the gradient would increase the height, and travel in the opposite direction would decrease the height. The minimal value of S would be attained when no local travel could decrease the function value, which is known as stationarity condition, stated as $\nabla S = 0$. Applying this to determining the coefficients (a_0, a_1) of a linear regression leads to the equations

$$\frac{\partial S}{\partial a_0} = 0 \Rightarrow 2 \sum_{i=1}^m (a_0 + a_1 x_i - y_i) = 0 \Leftrightarrow m a_0 + \left(\sum_{i=1}^m x_i \right) a_1 = \sum_{i=1}^m y_i,$$

$$\frac{\partial S}{\partial a_1} = 0 \Rightarrow 2 \sum_{i=1}^m (a_0 + a_1 x_i - y_i) x_i = 0 \Leftrightarrow \left(\sum_{i=1}^m x_i \right) a_0 + \left(\sum_{i=1}^m x_i^2 \right) a_1 = \sum_{i=1}^m x_i y_i.$$

The above calculations can become tedious, and do not illuminate the geometrical essence of the calculation, which can be brought out by reformulation in terms of a matrix-vector product that highlights the particular linear combination that is sought in a linear regression. Form a vector of errors with components $e_i = y(x_i) - y_i$, which for linear regression is $y(x) = a_0 + a_1 x$. Recognize that $y(x_i)$ is a linear combination of 1 and x_i with coefficients a_0, a_1 , or in vector form

$$\mathbf{e} = \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_m \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \end{pmatrix} - \mathbf{y} = (\mathbf{1} \ \mathbf{x}) \mathbf{a} - \mathbf{y} = \mathbf{A} \mathbf{a} - \mathbf{y}$$

The norm of the error vector $\|\mathbf{e}\|$ is smallest when $\mathbf{A} \mathbf{a}$ is as close as possible to \mathbf{y} . Since $\mathbf{A} \mathbf{a}$ is within the column space of $C(\mathbf{A})$, $\mathbf{A} \mathbf{a} \in C(\mathbf{A})$, the required condition is for \mathbf{e} to be orthogonal to the column space

$$\mathbf{e} \perp C(\mathbf{A}) \Rightarrow \mathbf{A}^T \mathbf{e} = \begin{pmatrix} \mathbf{1}^T \\ \mathbf{x}^T \end{pmatrix} \mathbf{e} = \begin{pmatrix} \mathbf{1}^T \mathbf{e} \\ \mathbf{x}^T \mathbf{e} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} = \mathbf{0}$$

$$\mathbf{A}^T \mathbf{e} = \mathbf{0} \Leftrightarrow \mathbf{A}^T (\mathbf{A} \mathbf{a} - \mathbf{y}) = \mathbf{0} \Leftrightarrow (\mathbf{A}^T \mathbf{A}) \mathbf{a} = \mathbf{A}^T \mathbf{y} = \mathbf{b}.$$

The above is known as the normal system, with $\mathbf{N} = \mathbf{A}^T \mathbf{A}$ is the normal matrix. The system $\mathbf{N} \mathbf{a} = \mathbf{b}$ can be interpreted as seeking the coordinates in the $\mathbf{N} = \mathbf{A}^T \mathbf{A}$ basis of the vector $\mathbf{b} = \mathbf{A}^T \mathbf{y}$. An example can be constructed by randomly perturbing a known function $y(x) = a_0 + a_1 x$ to simulate measurement noise and compare to the approximate $\tilde{\mathbf{a}}$ obtained by solving the normal system.

1. Generate some data on a line and perturb it by some random quantities

```
∴ m=100; x=LinRange(0,1,m); a=[2; 3];
∴ a0=a[1]; a1=a[2]; yex=a0 .+ a1*x; y=(yex+rand(m,1) .* 0.5);
∴
```

2. Form the matrices \mathbf{A} , $\mathbf{N} = \mathbf{A}^T \mathbf{A}$, vector $\mathbf{b} = \mathbf{A}^T \mathbf{y}$

```
∴ A=ones(m,2); A[:,2]=x; N=A' * A; b=A' * y;
∴
```

3. Solve the system $\mathbf{N} \mathbf{a} = \mathbf{b}$, and form the linear combination $\tilde{\mathbf{y}} = \mathbf{A} \mathbf{a}$ closest to \mathbf{y}

```
∴ atilde=N\b; [a atilde]
```

$$\begin{bmatrix} 2.0 & 1.9215699010834906 \\ 3.0 & 3.03714411616737 \end{bmatrix} \quad (6)$$

```
∴
```

The normal matrix basis $N = A^T A$ can however be an ill-advised choice. Consider $A \in \mathbb{R}^{2 \times 2}$ given by

$$A = [a_1 \ a_2] = \begin{bmatrix} 1 & \cos \theta \\ 0 & \sin \theta \end{bmatrix},$$

where the first column vector is taken from the identity matrix $a_1 = e_1$, and second is the one obtained by rotating it with angle θ . If $\theta = \pi/2$, the normal matrix is orthogonal, $A^T A = I$, but for small θ , A and $N = A^T A$ are approximated as

$$A \cong \begin{bmatrix} 1 & 1 \\ 0 & \theta \end{bmatrix}, N = [n_1 \ n_2] = \begin{bmatrix} 1 & 1 \\ 0 & \theta^2 \end{bmatrix}.$$

When θ is small a_1, a_2 are almost colinear, and n_1, n_2 even more so. This can lead to amplification of small errors, but can be avoided by recognizing that the best approximation in the 2-norm is identical to the Euclidean concept of orthogonal projection. The orthogonal projector onto $C(A)$ is readily found by QR -factorization, and the steps to solve least squares become

1. Compute $QR = A$
2. The projection of y onto the column space of A is $z = QQ^T y$, and has coordinates $c = Q^T y$ in the orthogonal basis Q .
3. The same z can also be obtained by linear combination of the columns of A , $z = Aa = QQ^T y$, and replacing A with its QR -factorization gives $QRa = Qc$, that leads to the system $Ra = c$, solved by back-substitution.

```
∴ Q,R=mgs(A); c=Q'*y;
```

```
∴ aQR=R\c; [a atilde aQR]
```

$$\begin{bmatrix} 2.0 & 1.9215699010834906 & 1.9065620791027633 \\ 3.0 & 3.03714411616737 & 3.0503545613518166 \end{bmatrix} \quad (7)$$

```
∴
```

The above procedure carried over to approximation by higher degree polynomials.

```
∴ m=100; n=6; x=LinRange(0,1,m); a=rand(-10:10,n,1); A=ones(m,1);
```

```
∴ for j=1:n-1
    global A
    A = [A x.^j];
end
```

```
∴ yex=A*a; y=yex .+ 0.001*(rand(m,1) .- 0.5); N=A'*A;
```

```
∴ b=A'*y;
```

```
∴ atilde=N\b; Q,R=mgs(A); c=Q'*y;
```

```
∴ aQR=R\c; [a atilde aQR]
```

$$\begin{bmatrix} 10.0 & 10.000017230385858 & 10.000017230388655 \\ -1.0 & -0.9992383469668031 & -0.9992383470406295 \\ 5.0 & 4.989621421635248 & 4.989621422094993 \\ -1.0 & -0.9668390169144284 & -0.9668390180268995 \\ -3.0 & -3.0392026170942916 & -3.0392026159405616 \\ -7.0 & -6.984207346901643 & -6.984207347332273 \end{bmatrix} \quad (8)$$

```
∴
```



Given data b , form A , find x , such that $\|e\| = \|Ax - b\|$ is minimized

$$e = b - Ax$$