CHAPTER 1

LEAST SQUARES

DATA COMPRESSION

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 $D = \{(x_i, y_i), i = 1, \ldots, m\}$. 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 $\boldsymbol{V} = [\ \boldsymbol{v}_1\ \ldots\ \boldsymbol{v}_n\ ] \in \mathbb{R}^{m \times n}$ is a basis for V and $\boldsymbol{W} = [\ \boldsymbol{w}_1\ \ldots\ \boldsymbol{w}_{m-n}\ ] \in \mathbb{R}^{m \times (m-n)}$ is a basis for W, then $\boldsymbol{U} = [\ \boldsymbol{v}_1\ \ldots\ \boldsymbol{v}_n\ \boldsymbol{w}_1\ \ldots\ \boldsymbol{w}_{m-n}\ ]$ is a basis for U. Even though the matrices $\boldsymbol{V}, \boldsymbol{W}$ are not necessarily square, they are said to be orthogonal, in the sense that all columns are of unit norm and orthogonal to one another. Computation of the matrix product $\boldsymbol{V}^T\boldsymbol{V}$ leads to the formation of the identity matrix within $\mathbb{R}^n$

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

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

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

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

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

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

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

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

```
octave] I4=eye(5); V=I4(:,1:2); P=V*V';␣Q=I4-P;
        b=rand(5,1); r=P*b; s=Q*b; disp([P b r s])

    1.00000    0.00000    0.00000    0.00000    0.00000    0.42253    0.42253    0.00000
    0.00000    1.00000    0.00000    0.00000    0.00000    0.95900    0.95900    0.00000
```

```
      0.00000    0.00000    0.00000    0.00000    0.00000    0.41781    0.00000    0.41781
      0.00000    0.00000    0.00000    0.00000    0.00000    0.45744    0.00000    0.45744
      0.00000    0.00000    0.00000    0.00000    0.00000    0.49784    0.00000    0.49784
octave]
```


Figure 1.1. Projection in $\mathbb{R}^2$. The vectors $r, s \in \mathbb{R}^2$ have two components, but could be expressed through scaling of e_1, e_2.

Returning to the general case, the orthogonal matrices $U \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{m \times n}$, $W \in \mathbb{R}^{m \times (m-n)}$ are associated with linear mappings $b = f(x) = Ux$, $r = g(b) = Pb$, $s = h(b) = (I - P)b$. The mapping f gives the components in the I basis of a vector whose components in the U basis are x. The mappings g, h project a vector onto span$\{v_1, \ldots, v_n\}$, span$\{w_1, \ldots, w_{m-n}\}$, respectively. When V, W are orthogonal matrices the projections are also orthogonal $r \perp s$. 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* $f \circ f = f$, *or if for any* $u \in U$, $f(f(u)) = f(u)$. *With* P *the matrix associated* f, *a projection matrix satisfies* $P^2 = P$.

$$P = VV^T$$

$$P^2 = PP = VV^TVV^T = V(V^TV)V^T = VIV^T = VV^T = P$$

2. Gram-Schmidt

Orthonormal vector sets $\{q_1, \ldots, q_n\}$ 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 $\{a_1, \ldots, a_n\}$. This is possible for independent vectors, through what is known as the Gram-Schmidt algorithm

1. Start with an arbitrary direction a_1

2. Divide by its norm to obtain a unit-norm vector $q_1 = a_1 / \|a_1\|$

3. Choose another direction a_2

4. Subtract off its component along previous direction(s) $a_2 - (q_1^T a_2) q_1$

5. Divide by norm $q_2 = (a_2 - (q_1^T a_2) q_1) / \|a_2 - (q_1^T a_2) q_1\|$

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

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

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 Octave.

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

```octave
octave] function [Q,R] = mgs(A)
          [m,n]=size(A); Q=A; R=eye(n);
          for i=1:n
            R(i,i) = sqrt(Q(:,i)'*Q(:,i));
            if (R(i,i)<eps) break;
            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;
        end
octave]
```

Note that the normalization condition $\| q_{ii} \| = 1$ is satisifed 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 applications.

```octave
octave] A=rand(4); [Q,R]=mgs(A); disp([Q R])
```

```
     0.82757   -0.25921   -0.49326    0.06802    0.83553    0.64827    1.24651    1.05301
     0.19408    0.53127    0.15805    0.80939    0.00000    0.93177    0.82700    0.87551
     0.22006    0.79553   -0.12477   -0.55058    0.00000    0.00000    0.38433   -0.20336
     0.47857   -0.13302    0.84625   -0.19270    0.00000    0.00000    0.00000    0.42469
```

```octave
octave] [Q1,R1]=qr(A); disp([Q1 R1])
```

```
    -0.82757    0.25921   -0.49326   -0.06802   -0.83553   -0.64827   -1.24651   -1.05301
    -0.19408   -0.53127    0.15805   -0.80939    0.00000   -0.93177   -0.82700   -0.87551
    -0.22006   -0.79553   -0.12477    0.55058    0.00000    0.00000    0.38433   -0.20336
    -0.47857    0.13302    0.84625    0.19270    0.00000    0.00000    0.00000   -0.42469
```

```octave
octave] disp([norm(A-Q*R) norm(A-Q1*R1)])
```

```
     1.1102e-16    8.0390e-16
```

```octave
octave]
```

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

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

$$AX = B, A\begin{bmatrix} x_1 & \dots & x_n \end{bmatrix} = \begin{bmatrix} Ax_1 & \dots & Ax_n \end{bmatrix}.$$

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

3. QR solution of linear algebra problems

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 $r_{ii} < \epsilon$ break;
> $x_i = c_i / r_{ii}$
>   for $j = i\text{-}1$ down to 1
> $c_j = c_j - r_{ji} x_i$
>   end
> end
> return x

```
octave] function x=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;
        end
octave]
```

The above operations are carried out in the background by the Octave backslash operation A\b to solve A*x=b, inspired by the scalar mnemonic $ax = b \Rightarrow x = (1/a)\,b$. Again, many additional refinements of the basic algorithm argue for using the built-in Octave functions, even though the above implementations can be verified as correct.

```
octave] xex=rand(4,1); b=A*xex; [Q,R]=mgs(A); c=Q'*b;␣x=bcks(R,c);␣x0=A\b;
octave]␣disp([xex␣x␣x0])

   ␣␣␣0.96838␣␣␣0.96838␣␣␣0.96838
   ␣␣␣0.31829␣␣␣0.31829␣␣␣0.31829
   ␣␣␣0.58529␣␣␣0.58529␣␣␣0.58529
   ␣␣␣0.38250␣␣␣0.38250␣␣␣0.38250

octave]␣
```

3.2. General orthogonal bases

The above approch for the real vector space $\mathscr{R}_m$ 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 $\mathscr{C}^\infty[-1, 1]$ that can be differentiated an arbitrary number of times, the Taylor series

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

is seen to be a linear combination of the monomial basis $M = \begin{bmatrix} 1 & x & x^2 & \ldots \end{bmatrix}$ with scaling coefficients $\{f(0), f'(0), \tfrac{1}{2} f''(0), \ldots\}$. 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

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

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 $x_i = (i - 1)\,h - 1$, $h = 2/(m - 1)$, $i = 1, \ldots, m$, by approximation of the integral by a Riemann sum

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

```
octave]␣m=50;␣h=2/(m-1);␣x=(-1:h:1)'; M=[x.^0 x.^1 x.^2 x.^3 x.^4]; [Q,R]=mgs(M);
        S=diag(1./Q(m,:)); P=Q*S; sc=[-1 1 -1 1];
        figure(1); plot(x,M(:,1),x,M(:,2),x,M(:,3),x,M(:,4)); axis(sc); grid on;
        figure(2); plot(x,P(:,1),x,P(:,2),x,P(:,3),x,P(:,4)); axis(sc); grid on;
```
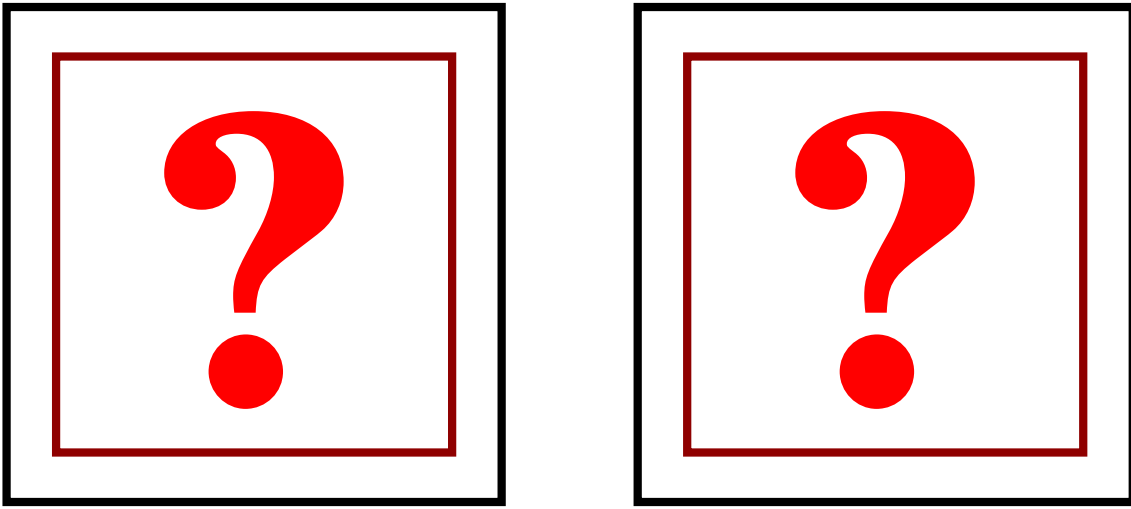
octave]

Figure 1.2. 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 $D = \{(x_i, y_i) \mid i = 1, \ldots, m\}$ 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 $S(a_0, a_1) \geqslant 0$ 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

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

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

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 liner 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

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

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

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

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

The above is known as the normal system, with $N = A^T A$ is the normal matrix. The system $Na = b$ can be interpreted as seeking the coordinates in the $N = A^T A$ basis of the vector $b = A^T y$. 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 $\tilde{a}$ obtained by solving the normal system.

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

```
octave] m=100; x=(0:m-1)/m; a=[2; 3];
        a0=a(1); a1=a(2); yex=a0+a1*x; y=(yex+rand(1,m)-0.5)';
octave]
```

2. Form the matrices A, $N = A^T A$, vector $b = A^T y$

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

3. Solve the system $Na = b$, and form the linear combination $\tilde{y} = Aa$ closest to y

```
octave] atilde=N\b; disp([a atilde]);
        2.0000    2.0302
        3.0000    2.9628
octave]
```

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

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

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

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

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 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.

```
octave] [Q,R]=qr(A); c=Q'*y; aQR=R\c; disp([a atilde aQR])
        2.0000    2.0302    2.0302
        3.0000    2.9628    2.9628
octave]
```

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

```octave
octave] m=100; n=6; x=(0:m-1)/m; x=x';␣a=randi(10,n,1);␣A=[];
        for j=1:n
          A = [A x.^(j-1)];
        end;
        yex=A*a; y=yex+(rand(m,1)-0.5);
octave] N=A'*A;␣b=A'*y; atilde=inv(N)*b;
        [Q,R]=qr(A); c=Q'*y;␣aQR=R\c;␣
        disp([a atilde aQR]);
         8.0000    8.0847    8.0847
         8.0000    7.1480    7.1480
         4.0000    4.2264    4.2264
         4.0000    8.7568    8.7568
        10.0000    2.7420    2.7420
         6.0000    9.0386    9.0386

octave]
```


Given data y, form A, find a, such that $\|e\| = \|Aa - y\|$ is minimized

MODEL REDUCTION

1. Projection of mappings

The least-squares problem

$$\min_{x \in \mathbb{R}^n} \| y - Ax \| \tag{1.1}$$

focuses on a simpler representation of a data vector $y \in \mathbb{R}^m$ as a linear combination of column vectors of $A \in \mathbb{R}^{m \times n}$. Consider some phenomenon modeled as a function between vector spaces $f : X \to Y$, such that for input parameters $x \in X$, the state of the system is $y = f(x)$. For most models f is differentiable, a transcription of the condition that the system should not exhibit jumps in behavior when changing the input parameters. Then by appropriate choice of units and origin, a linearized model

$$y = Ax, A \in \mathbb{R}^{m \times n},$$

is obtained if $y \in C(A)$, expressed as (1) if $y \notin C(A)$.

A simpler description is often sought, typically based on recognition that the inputs and outputs of the model can themselves be obtained as linear combinations $x = Bu$, $y = Cv$, involving a smaller set of parameters $u \in \mathbb{R}^q$, $v \in \mathbb{R}^p$, $p < m$, $q < n$. The column spaces of the matrices $B \in \mathbb{R}^{n \times q}$, $C \in \mathbb{R}^{m \times p}$ are vector subspaces of the original set of inputs and outputs, $C(B) \leq \mathbb{R}^n$, $C(C) \leq \mathbb{R}^m$. The sets of column vectors of B, C each form a *reduced basis* for the system inputs and outputs if they are chosed to be of full rank. The reduced bases are assumed to have been orthonormalized through the Gram-Schmidt procedure such that $B^T B = I_q$, and $C^T C = I_p$. Expressing the model inputs and outputs in terms of the reduced basis leads to

$$Cv = ABu \Longrightarrow v = C^T ABu \Longrightarrow v = Ru.$$

The matrix $R = C^T AB$ is called the *reduced system matrix* and is associated with a mapping $g: U \rightarrow V$, that is a restriction to the U, V vector subspaces of the mapping f. When f is an endomorphism, $f: X \rightarrow X$, $m = n$, the same reduced basis is used for both inputs and outputs, $x = Bu$, $y = Bv$, and the reduced system is

$$v = Ru, R = B^T AB.$$

Since B is assumed to be orthogonal, the projector onto $C(B)$ is $P_B = BB^T$. Applying the projector on the inital model

$$P_B y = P_B Ax$$

leads to $BB^T y = BB^T Ax$, and since $v = B^T y$ the relation $Bv = BB^T ABu$ is obtained, and conveniently grouped as

$$Bv = B(B^T AB)u \Longrightarrow Bv = B(Ru),$$

again leading to the reduced model $v = Bu$. The above calculation highlights that the reduced model is a projection of the full model $y = Ax$ on $C(B)$.

2. Reduced bases

2.1. Correlation matrices

Correlation coefficient. Consider two functions $x_1, x_2: \mathbb{R} \rightarrow \mathbb{R}$, that represent data streams in time of inputs $x_1(t)$ and outputs $x_2(t)$ of some system. A basic question arising in modeling and data science is whether the inputs and outputs are themselves in a functional relationship. This usually is a consequence of incomplete knowledge of the system, such that while x_1, x_2 might be assumed to be the most relevant input, output quantities, this is not yet fully established. A typical approach is to then carry out repeated measurements leading to a data set $D = \{(x_1(t_i), x_2(t_i)) \mid i = 1,...,N\}$, thus defining a relation. Let $x_1, x_2 \in \mathbb{R}^N$ denote vectors containing the input and output values. The *mean values* μ_1, μ_2 of the input and output are estimated by the statistics

$$\mu_1 \cong \bar{x}_1 = \frac{1}{N}\sum_{i=1}^{N} x_1(t_i) = E[x_1], \mu_2 \cong \bar{x}_2 = \frac{1}{N}\sum_{i=1}^{N} x_2(t_i) = E[x_2],$$

where E is the expectation seen to be a linear mapping, $E: \mathbb{R}^N \rightarrow \mathbb{R}$ whose associated matrix is

$$E = \frac{1}{N}\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix},$$

and the means are also obtained by matrix vector multiplication (linear combination),

$$\bar{x}_1 = Ex_1, \bar{x}_2 = Ex_2.$$

Deviation from the mean is measured by the *standard deviation* defined for x_1, x_2 by

$$\sigma_1 = \sqrt{E[(x_1 - \mu_1)^2]}, \sigma_2 = \sqrt{E[(x_2 - \mu_2)^2]}.$$

Note that the standard deviations are no longer linear mappings of the data.

Assume that the origin is chosen such that $\bar{x}_1 = \bar{x}_2 = 0$. One tool to estalish whether the relation D is also a function is to compute the *correlation coefficient*

$$\rho(x_1, x_2) = \frac{E[x_1 x_2]}{\sigma_1 \sigma_2} = \frac{E[x_1 x_2]}{\sqrt{E[x_1^2] E[x_2^2]}},$$

that can be expressed in terms of a scalar product and 2-norm as

$$\rho(x_1, x_2) = \frac{x_1^T x_2}{\|x_1\| \|x_2\|}.$$

Squaring each side of the norm property $\|x_1 + x_2\| \leq \|x_1\| + \|x_2\|$, leads to

$$(x_1 + x_2)^T (x_1 + x_2) \leq x_1^T x_1 + x_2^T x_2 + 2 \|x_1\| \|x_2\| \implies x_1^T x_2 \leq \|x_1\| \|x_2\|,$$

known as the Cauchy-Schwarz inequality, which implies $-1 \leq \rho(x_1, x_2) \leq 1$. Depending on the value of ρ, the variables $x_1(t), x_2(t)$ are said to be:

1. *uncorrelated*, if $\rho = 0$;

2. *correlated*, if $\rho = 1$;

3. *anti-correlated*, if $\rho = -1$.

The numerator of the correlation coefficient is known as the covariance of x_1, x_2

$$\text{cov}(x_1, x_2) = E[x_1 x_2].$$

The correlation coefficient can be interpreted as a normalization of the covariance, and the relation

$$\text{cov}(x_1, x_2) = x_1^T x_2 = \rho(x_1, x_2) \|x_1\| \|x_2\|,$$

is the two-variable version of a more general relationship encountered when the system inputs and outputs become vectors.

Patterns in data. Consider now a related problem, whether the input and output parameters $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$ thought to characterize a system are actually well chosen, or whether they are redundant in the sense that a more insightful description is furnished by $u \in \mathbb{R}^q$, $v \in \mathbb{R}^p$ with fewer components $p < m, q < n$. Applying the same ideas as in the correlation coefficient, a sequence of N measurements is made leading to data sets

$$X = [\ x_1\ \ x_2\ \ \ldots\ \ x_n\ ] \in \mathbb{R}^{N \times n}, Y = [\ y_1\ \ y_2\ \ \ldots\ \ y_n\ ] \in \mathbb{R}^{N \times m}.$$

Again, by appropriate choice of the origin the means of the above measurements is assumed to be zero

$$E[x] = 0, E[y] = 0.$$

Covariance matrices can be constructed by

$$C_X = X^T X = \begin{bmatrix} x_1^T \\ x_2^T \\ \vdots \\ x_n^T \end{bmatrix} [\ x_1\ \ x_2\ \ \ldots\ \ x_n\ ] = \begin{bmatrix} x_1^T x_1 & x_1^T x_2 & \ldots & x_1^T x_n \\ x_2^T x_1 & x_2^T x_2 & \ldots & x_2^T x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_n^T x_1 & x_n^T x_2 & \ldots & x_n^T x_n \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

Consider now the SVDs of $C_X = N \Lambda N^T$, $X = U \Sigma S^T$, and from

$$C_X = X^T X = (U \Sigma S^T)^T U \Sigma S^T = S \Sigma^T U^T U \Sigma S^T = S \Sigma^T \Sigma S^T = N \Lambda N^T,$$

identify $N = S$, and $\Lambda = \Sigma^T \Sigma$.

Recall that the SVD returns an order set of singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq$, and associated singular vectors. In many applications the singular values decrease quickly, often exponentially fast. Taking the first q singular modes then gives a basis set suitable for mode reduction

$$x = S_q u = [\ s_1\ \ s_2\ \ \ldots\ \ s_q\ ] u.$$