Properties of Matrices and Operations on Matrices

A very useful factorization is

$$A = QR,$$

where Q is orthogonal and R is upper triangular or trapezoidal.

This is called the QR factorization.

Forms of the Factors

If A is square and of full rank, R has the form

$$\begin{bmatrix} \mathtt{X} & \mathtt{X} & \mathtt{X} \\ \mathtt{0} & \mathtt{X} & \mathtt{X} \\ \mathtt{0} & \mathtt{0} & \mathtt{X} \end{bmatrix}.$$

If A is nonsquare, R is nonsquare, with an upper triangular submatrix.

If A has more columns than rows, R is trapezoidal and can be written as $[R_1 \,|\, R_2]$, where R_1 is upper triangular.

If A is $n \times m$ with more rows than columns, which is the case in common applications of QR factorization, then

$$R = \begin{bmatrix} R_1 \\ 0 \end{bmatrix},$$

where R_1 is $m \times m$ upper triangular.

When A has more rows than columns, we can likewise partition Q as $[Q_1 \,|\, Q_2]$, and we can use a version of Q that contains only relevant rows or columns,

$$A = Q_1 R_1,$$

where Q_1 is an $n \times m$ matrix whose columns are orthonormal.

This form is called a "skinny" QR.

It is more commonly used than one with a square Q.

Relation to the Moore-Penrose Inverse

It is interesting to note that the Moore-Penrose inverse of A with full column rank is immediately available from the QR factorization:

$$A^{+} = \left[\begin{array}{cc} R_1^{-1} & 0 \end{array} \right] Q^{\mathsf{T}}.$$

Nonfull Rank Matrices

If A is square but not of full rank, R has the form

$$\begin{bmatrix} \mathtt{X} & \mathtt{X} & \mathtt{X} \\ \mathtt{0} & \mathtt{X} & \mathtt{X} \\ \mathtt{0} & \mathtt{0} & \mathtt{0} \end{bmatrix}.$$

In the common case in which A has more rows than columns, if A is not of full (column) rank, R_1 will have the form shown above.

If A is not of full rank, we apply permutations to the columns of A by multiplying on the right by a permutation matrix.

The permutations can be taken out by a second multiplication on the right.

If A is of rank r ($\leq m$), the resulting decomposition consists of three matrices: an orthogonal Q, a T with an $r \times r$ upper triangular submatrix, and a permutation matrix E_π^T,

$$A = QTE_\pi^\mathsf{T}.$$

The matrix T has the form

$$T = \begin{bmatrix} T_1 & T_2 \\ 0 & 0 \end{bmatrix},$$

where T_1 is upper triangular and is $r \times r$.

The decomposition is not unique because of the permutation matrix.

The choice of the permutation matrix is the same as the pivoting that we discussed in connection with Gaussian elimination.

A generalized inverse of A is immediately available:

$$A^{-} = P \begin{bmatrix} T_1^{-1} & 0 \\ 0 & 0 \end{bmatrix} Q^{\mathsf{T}}.$$

Additional orthogonal transformations can be applied from the right-hand side of the $n \times m$ matrix A to yield

$$A = QRU^{\mathsf{T}},$$

where R has the form

$$R = \begin{bmatrix} R_1 & 0 \\ 0 & 0 \end{bmatrix},$$

where R_1 is $r \times r$ upper triangular, Q is $n \times n$, and U^{T} is $n \times m$ and orthogonal.

The decomposition is unique, and it provides the unique Moore-Penrose generalized inverse of A:

$$A^{+} = U \begin{bmatrix} R_1^{-1} & 0 \\ 0 & 0 \end{bmatrix} Q^{\mathsf{T}}.$$

It is often of interest to know the rank of a matrix.

Given a decomposition of this form, the rank is obvious, and in practice, this QR decomposition with pivoting is a good way to determine the rank of a matrix.

The QR decomposition is said to be “rank-revealing”.

The computations are quite sensitive to rounding, however, and the pivoting must be done with some care.

The QR factorization is particularly useful in computations for overdetermined systems, and in other computations involving nonsquare matrices.

Formation of the QR Factorization

There are three good methods for obtaining the QR factorization: Householder transformations or reflections; Givens transformations or rotations; and the (modified) Gram-Schmidt procedure.

Different situations may make one of these procedures better than the two others.

The Householder transformations are probably the most commonly used.

If the data are available only one row at a time, the Givens transformations are very convenient.

Whichever method is used to compute the QR decomposition, at least $2n^3/3$ multiplications and additions are required. The operation count is therefore about twice as great as that for an LU decomposition.

Reflections Added after Lecture

Rotations are special geometric transformations discussed in the lecture on September 24 that preserve the Euclidean length of a vector.

Reflections are a special kind of rotation.

Let u and v be orthonormal vectors, and let x be a vector in the space spanned by u and v, so

$$x = c_1 u + c_2 v$$

for some scalars c_1 and c_2.

The vector

$$\tilde{x} = -c_1 u + c_2 v$$

is a *reflection* of x through the line defined by the vector v, or $u^\perp$.

This reflection is a rotation in the plane defined by u and v through an angle of twice the size of the angle between x and v.

The form of $\tilde{x}$ of course depends on the vector v and its relationship to x. In a common application of reflections in linear algebraic computations, we wish to rotate a given vector into a vector collinear with a coordinate axis; that is, we seek a reflection that transforms a vector

$$x = (x_1, x_2, \ldots, x_n)$$

into a vector collinear with a unit vector,

$$\begin{aligned} \tilde{x} &= (0, \ldots, 0, \tilde{x}_i, 0, \ldots, 0) \\ &= \pm\|x\|_2 e_i. \end{aligned}$$

Geometrically, in two dimensions we have the picture shown below, where $i = 1$.

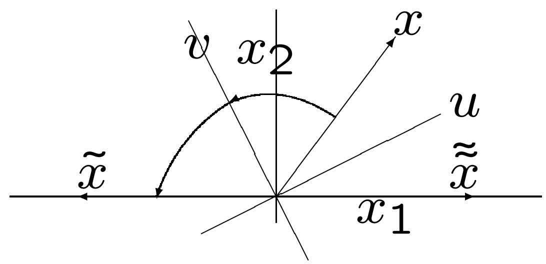

Which vector x is rotated through (that is, which is u and which is v) depends on the choice of the sign in $\pm\|x\|_2$.

The choice that was made yields the $\tilde{x}$ shown in the figure, and from the figure, this can be seen to be correct.

Note that

$$v = \frac{1}{|2c_2|}(x + \tilde{x})$$

If the opposite choice is made, we get the $\tilde{\tilde{x}}$ shown. In the simple two-dimensional case, this is equivalent to reversing our choice of u and v.

Householder Reflections

Consider the problem of reflecting x through the vector v. As before, we assume that u and v are orthonormal vectors and that x lies in a space spanned by u and v, and $x = c_1 u + c_2 v$. Form the matrix

$$H = I - 2uu^{\mathsf{T}},$$

and note that

$$\begin{aligned} Hx &= c_1 u + c_2 v - 2c_1 uu^{\mathsf{T}}u - 2c_2 uu^{\mathsf{T}}v \\ &= c_1 u + c_2 v - 2c_1 u^{\mathsf{T}}uu - 2c_2 u^{\mathsf{T}}vu \\ &= -c_1 u + c_2 v \\ &= \tilde{x}. \end{aligned}$$

The matrix H is a *reflector*; it has transformed x into its reflection $\tilde{x}$ about v.

A reflection is also called a Householder reflection or a Householder transformation, and the matrix H is called a Householder matrix or a Householder reflector.

The following properties of H are immediate:

- $Hu = -u$.

- $Hv = v$ for any v orthogonal to u.

- $H = H^{\mathsf{T}}$ (symmetric).

- $H^{\mathsf{T}} = H^{-1}$ (orthogonal).

Because H is orthogonal, if $Hx = \tilde{x}$, then $\|x\|_2 = \|\tilde{x}\|_2$, so $\tilde{x}_1 = \pm\|x\|_2$.

The matrix uu^{T} is symmetric, idempotent, and of rank 1.

A transformation by a matrix of the form $A - vw^{\mathrm{T}}$ is often called a "rank-one" update, because vw^{T} is of rank 1.

Thus, a Householder reflection is a special rank-one update.

Zeroing Elements in a Vector

The usefulness of Householder reflections results from the fact that it is easy to construct a reflection that will transform a vector x into a vector $\tilde{x}$ that has zeros in all but one position.

To construct the reflector of x into $\tilde{x}$, we first form the vector about which to reflect x. The vector about which we perform the reflection is merely

$$x + \tilde{x}.$$

Because $\|\tilde{x}\|_2 = \|x\|_2$, we know $\tilde{x}$ to within the sign; that is,

$$\tilde{x} = (0, \ldots, 0, \pm\|x\|_2, 0, \ldots, 0).$$

We choose the sign so as not to add quantities of different signs and possibly similar magnitudes.

Hence, we have

$$q = (x_1,\ \ldots,\ x_{i-1},\ x_i + \mathrm{sign}(x_i)\|x\|_2,\ x_{i+1},\ \ldots,\ x_n),$$

then

$$u = q/\|q\|_2,$$

and finally

$$H = I - 2uu^{\mathsf{T}}.$$

Consider, for example, the vector

$$x = (3, 1, 2, 1, 1),$$

which we wish to transform into

$$\tilde{x} = (\widetilde{x_1}, 0, 0, 0, 0).$$

We have

$$\|x\| = 4,$$

so we form the vector

$$u = \frac{1}{\sqrt{56}}(7, 1, 2, 1, 1).$$

So

$$
\begin{aligned}
H &= I - 2uu^{\mathsf{T}} \\
&= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} - \frac{1}{28} \begin{bmatrix} 49 & 7 & 14 & 7 & 7 \\ 7 & 1 & 2 & 1 & 1 \\ 14 & 2 & 4 & 2 & 2 \\ 7 & 1 & 2 & 1 & 1 \\ 7 & 1 & 2 & 1 & 1 \end{bmatrix} \\
&= \frac{1}{28} \begin{bmatrix} -21 & -7 & -14 & -7 & -7 \\ -7 & 27 & -2 & -1 & -1 \\ -14 & -2 & 24 & -2 & -2 \\ -7 & -1 & -2 & 27 & -1 \\ -7 & -1 & -2 & -1 & 27 \end{bmatrix}
\end{aligned}
$$

to yield $Hx = (-4, 0, 0, 0, 0)$.

Householder Reflections to Form the QR Factorization

To use reflectors to compute a QR factorization, we form in sequence the reflector for the i^{th} column that will produce 0s below the (i, i) element.

For a convenient example, consider the matrix

$$A = \begin{bmatrix} 3 & -\frac{98}{28} & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 1 & \frac{122}{28} & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 2 & -\frac{8}{28} & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 1 & \frac{66}{28} & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 1 & \frac{10}{28} & \mathtt{X} & \mathtt{X} & \mathtt{X} \end{bmatrix}.$$

The first transformation would be determined so as to transform $(3, 1, 2, 1, 1)$ to $(\mathtt{X}, 0, 0, 0, 0)$.

Call this first Householder matrix P_1. We have

$$P_1 A = \begin{bmatrix} -4 & 1 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 5 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 1 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 3 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 1 & \mathtt{X} & \mathtt{X} & \mathtt{X} \end{bmatrix}.$$

We now choose a reflector to transform $(5, 1, 3, 1)$ to $(-6, 0, 0, 0)$. We do not want to disturb the first column in $P_1 A$ shown above, so we form P_2 as

$$P_2 = \begin{bmatrix} 1 & 0 & \ldots & 0 \\ 0 & & & \\ \vdots & & H_2 & \\ 0 & & & \end{bmatrix}.$$

Forming the vector $(11, 1, 3, 1)/\sqrt{132}$ and proceeding as before, we get the reflector

$$\begin{aligned} H_2 &= I - \frac{1}{66}(11, 1, 3, 1)(11, 1, 3, 1)^{\mathsf{T}} \\ &= \frac{1}{66}\begin{bmatrix} -55 & -11 & -33 & -11 \\ -11 & 65 & -3 & -1 \\ -33 & -3 & 57 & -3 \\ -11 & -1 & -3 & 65 \end{bmatrix}. \end{aligned}$$

Now we have

$$P_2 P_1 A = \begin{bmatrix} -4 & 1 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & -6 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 0 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 0 & \mathtt{X} & \mathtt{X} & \mathtt{X} \\ 0 & 0 & \mathtt{X} & \mathtt{X} & \mathtt{X} \end{bmatrix}.$$

Continuing in this way for three more steps, we would have the QR decomposition of A with $Q^{\mathsf{T}} = P_5 P_4 P_3 P_2 P_1$.

The number of computations for the QR factorization of an $n \times n$ matrix using Householder reflectors is $2n^3/3$ multiplications and $2n^3/3$ additions.

Givens Rotations to Form the QR Factorization

Just as we built the QR factorization by applying a succession of Householder reflections, we can also apply a succession of Givens rotations to achieve the factorization.

If the Givens rotations are applied directly, the number of computations is about twice as many as for the Householder reflections, but if fast Givens rotations are used and accumulated cleverly, the number of computations for Givens rotations is not much greater than that for Householder reflections.

It is necessary to monitor the differences in the magnitudes of the elements in the C matrix and often necessary to rescale the elements.

This additional computational burden is excessive unless done carefully for a description of an efficient method).

Gram-Schmidt Transformations to Form the QR Factorization

Gram-Schmidt transformations yield a set of orthonormal vectors that span the same space as a given set of linearly independent vectors, $\{x_1, x_2, \ldots, x_m\}$.

Application of these transformations is called Gram-Schmidt orthogonalization.

If the given linearly independent vectors are the columns of a matrix A, the Gram-Schmidt transformations ultimately yield the QR factorization of A.

Applications in Regression and Other Linear Computations

We write and overdetermined system $Ab \approx y$ as

$$Xb = y - r,$$

where r is an n-vector of possibly arbitrary residuals or "errors".

A *least squares* solution $\widehat{b}$ to the system is one such that the Euclidean norm of the vector of residuals is minimized; that is, the solution to the problem

$$\min_b \|y - Xb\|_2.$$

The least squares solution is also called the "ordinary least squares" (OLS) fit.

By rewriting the square of this norm as

$$(y - Xb)^{\mathsf{T}}(y - Xb),$$

differentiating, and setting it equal to 0, we see that the minimum (of both the norm and its square) occurs at the $\widehat{b}$ that satisfies the square system

$$X^{\mathsf{T}}X\widehat{b} = X^{\mathsf{T}}y.$$

The system is called the *normal equations*.

Because the condition number of $X^{\mathsf{T}}X$ is the square of the condition number of X, it may be better to work directly on X rather than to use the normal equations.

The normal equations are useful expressions, however, whether or not they are used in the computations. This is another case where *a formula does not define an algorithm.*

Special Properties of Least Squares Solutions

The least squares fit to the overdetermined system has a very useful property with two important consequences. The least squares fit partitions the space into two interpretable orthogonal spaces. As we see from the equation, the residual vector $y - X\widehat{b}$ is orthogonal to each column in X:

$$X^{\mathsf{T}}(y - X\widehat{b}) = 0.$$

A consequence of this fact for models that include an intercept is that the sum of the residuals is 0. (The residual vector is orthogonal to the 1 vector.) Another consequence for models that include an intercept is that the least squares solution provides an exact fit to the mean.

These properties are so familiar to statisticians that some think that they are essential characteristics of any regression modeling; they are not.

Weighted Least Squares

One of the simplest variations on fitting the linear model $Xb \approx y$ is to allow different weights on the observations; that is, instead of each row of X and corresponding element of y contributing equally to the fit, the elements of X and y are possibly weighted differently. The relative weights can be put into an n-vector w and the squared norm replaced by a quadratic form in $\mathrm{diag}(w)$. More generally, we form the quadratic form as

$$(y - Xb)^{\mathsf{T}} W (y - Xb),$$

where W is a positive definite matrix. Because the weights apply to both y and Xb, there is no essential difference in the weighted or unweighted versions of the problem.

The use of the QR factorization for the overdetermined system in which the weighted norm is to be minimized is similar to the development above. It is exactly what we get if we replace $y - Xb$ by $W_{\mathsf{C}}(y - Xb)$, where W_{C} is the Cholesky factor of W.

Numerical Accuracy in Overdetermined Systems

Bounds on the numerical error can be expressed in terms of the condition number of the coefficient matrix, which is the ratio of norms of the coefficient matrix and its inverse.

One of the most useful versions of this condition number is the one using the L_2 matrix norm, which is called the spectral condition number.

This is the most commonly used condition number, and we generally just denote it by $\kappa(\cdot)$.

The spectral condition number is the ratio of the largest eigenvalue in absolute value to the smallest in absolute value, and this extends easily to a definition of the spectral condition number that applies both to nonsquare matrices and to singular matrices: the *condition number* of a matrix is the ratio of the largest singular value to the smallest nonzero singular value.

More on Condition Numbers

The nonzero singular values of X are the square roots of the nonzero eigenvalues of $X^\top X$; hence

$$\kappa(X^\top X) = (\kappa(X))^2.$$

The condition number of $X^\top X$ is a measure of the numerical accuracy we can expect in solving the normal equations.

Because the condition number of X is smaller, we have an indication that it might be better not to form the normal equations unless we must.

It might be better to work just with X.

Least Squares with a Full Rank Coefficient Matrix

If the $n \times m$ matrix X is of full column rank, the least squares solution, is $\widehat{b} = (X^\top X)^{-1} X^\top y$ and is obviously unique.

A good way to compute this is to form the QR factorization of X.

First we write $X = QR$ where R is

$$R = \begin{bmatrix} R_1 \\ 0 \end{bmatrix},$$

with R_1 an $m \times m$ upper triangular matrix.

Regression Computations

The residual norm can be written as

$$\begin{aligned}(y - Xb)^{\mathsf{T}}(y - Xb) &= (y - QRb)^{\mathsf{T}}(y - QRb) \\ &= (Q^{\mathsf{T}}y - Rb)^{\mathsf{T}}(Q^{\mathsf{T}}y - Rb) \\ &= (c_1 - R_1 b)^{\mathsf{T}}(c_1 - R_1 b) + c_2^{\mathsf{T}} c_2,\end{aligned}$$

where c_1 is a vector with m elements and c_2 is a vector with $n - m$ elements, such that

$$Q^{\mathsf{T}}y = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

Because quadratic forms are nonnegative, the minimum of the residual norm occurs when $(c_1 - R_1 b)^{\mathsf{T}}(c_1 - R_1 b) = 0$; that is, when $(c_1 - R_1 b) = 0$, or

$$R_1 b = c_1.$$

Because R_1 is triangular, the system is easy to solve: $\widehat{b} = R_1^{-1} c_1$.

We have

$$X^{+} = \left[\begin{array}{cc} R_1^{-1} & 0 \end{array}\right] Q^{\mathsf{T}},$$

and so we have

$$\widehat{b} = X^{+} y.$$

The minimum of the residual norm is $c_2^{\mathsf{T}} c_2$.

This is called the *residual sum of squares* in the least squares fit.

Least Squares with a Coefficient Matrix Not of Full Rank

If X is not of full rank (that is, if X has rank $r < m$), the least squares solution is not unique, and in fact a solution is any vector $\widehat{b} = (X^{\mathsf{T}}X)^{-}X^{\mathsf{T}}y$, where $(X^{\mathsf{T}}X)^{-}$ is any generalized inverse.

This is a solution to the normal equations.

The residual corresponding to this solution is

$$y - X(X^{\mathsf{T}}X)^{-}X^{\mathsf{T}}y = (I - X(X^{\mathsf{T}}X)^{-}X^{\mathsf{T}})y.$$

The residual vector is invariant to the choice of generalized inverse.

An Optimal Property of the Solution Using the Moore-Penrose Inverse

The solution corresponding to the Moore-Penrose inverse is unique because, as we have seen, that generalized inverse is unique.

That solution is interesting for another reason, however: the b from the Moore-Penrose inverse has the minimum L_2-norm of all solutions.

To see that this solution has minimum norm, first factor X,

$$X = QRU^{\mathsf{T}},$$

and form the Moore-Penrose inverse:

$$X^{+} = U \begin{bmatrix} R_1^{-1} & 0 \\ 0 & 0 \end{bmatrix} Q^{\mathsf{T}}.$$

So

$$\widehat{b} = X^{+} y$$

is a least squares solution, just as in the full rank case.

Now, let

$$Q^{\mathsf{T}} y = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$$

and let

$$U^{\mathsf{T}} b = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix},$$

where z_1 has r elements.

We seek to minimize $\|y - Xb\|_2$; and because multiplication by an orthogonal matrix does not change the norm, we have

$$\begin{aligned} \|y - Xb\|_2 &= \|Q^\top (y - XUU^\top b)\|_2 \\ &= \left\| \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} - \begin{bmatrix} R_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \right\|_2 \\ &= \left\| \begin{pmatrix} c_1 - R_1 z_1 \\ c_2 \end{pmatrix} \right\|_2 . \end{aligned}$$

The residual norm is minimized for $z_1 = R_1^{-1} c_1$ and z_2 arbitrary. However, if $z_2 = 0$, then $\|z\|_2$ is also minimized.

Because $U^\top b = z$ and U is orthogonal, $\|\widehat{b}\|_2 = \|z\|_2$, and so $\|\widehat{b}\|_2$ is the minimum among all least squares solutions.