

CS 4220: Numerical Analysis

Least squares and QR

David Bindel

2026-02-23

Orthogonal transformations and Gram-Schmidt

We saw in the last lecture that a natural decomposition for thinking about least squares problems is the QR decomposition

$$A = QR,$$

where Q is an $m \times m$ orthogonal matrix and R is an $m \times n$ upper triangular matrix. Equivalently, we can write the “economy” version of the decomposition, $A = QR$ with an $m \times n$ matrix Q and an $n \times n$ upper triangular R , where the columns of Q form an orthonormal basis for the range space of A . Using this decomposition, we can solve the least squares problem via the triangular system

$$Rx = Q^T b.$$

The *Gram-Schmidt* procedure is usually the first method people learn to convert some existing basis (columns of A) into an orthonormal basis (columns of Q). For each column of A , the procedure subtracts off any components in the direction of the previous columns, and then scales the remainder to be unit length. In Julia, Gram-Schmidt looks something like this:

```
function orth_cgs0(A)
    m,n = size(A)
    Q = zeros(m,n)
    for j = 1:n
        v = A[:,j] # Take the jth original basis vector
        v = v-Q[:,1:j-1]*(Q[:,1:j-1]'*v) # Orthogonalize vs q_1, ... q_{j-1}
        v = v/norm(v) # Normalize what remains
        Q[:,j] = v # Add result to Q basis
    end
    Q
end
```

Where does R appear in this algorithm? It appears thus:

```
function orth_cgs(A)
    m,n = size(A)
    Q = zeros(m,n)
    R = zeros(n,n)
    for j = 1:n
        v = A[:,j] # Take the jth original basis vector
        R[1:j-1,j] = Q[:,1:j-1]'*v # Project onto q_1, ..., q_{j-1}
        v = v-Q[:,1:j-1]*R[1:j-1,j] # Orthogonalize vs q_1, ... q_{j-1}
        R[j,j] = norm(v) # Compute normalization constant
        v = v/R[j,j] # Normalize what remains
        Q[:,j] = v # Add result to Q basis
    end
    Q, R
end
```

That is, R accumulates the multipliers that we computed from the Gram-Schmidt procedure. This idea that the multipliers in an algorithm can be thought of as entries in a matrix should be familiar, since we encountered it before when we looked at Gaussian elimination.

Householder transformations

The Gram-Schmidt orthogonalization procedure is not generally recommended for numerical use. Suppose we write $A = [a_1 \dots a_m]$ and $Q = [q_1 \dots q_m]$. The essential problem is that if $r_{jj} \ll \|a_j\|_2$, then cancellation can destroy the accuracy of the computed q_j ; and in particular, the computed q_j may not be particularly orthogonal to the previous q_j . Actually, loss of orthogonality can build up even if the diagonal elements of R are not exceptionally small. This is Not Good, and while we have some tricks to mitigate the problem, we need a different approach if we want the problem to go away.

Recall that one way of expressing the Gaussian elimination algorithm is in terms of Gauss transformations that serve to introduce zeros into the lower triangle of a matrix. *Householder* transformations are orthogonal transformations (reflections) that can be used to similar effect. Reflection across the plane orthogonal to a unit normal vector v can be expressed in matrix form as

$$H = I - 2vv^T.$$

Now suppose we are given a vector x and we want to find a reflection that transforms x into a direction parallel to some unit vector y . The right reflection is through a hyperplane that bisects

the angle between x and y (see Figure 1), which we can construct by taking the hyperplane normal to $x - \|x\|y$. That is, letting $u = x - \|x\|y$ and $v = u/\|u\|$, we have

$$\begin{aligned}(I - 2vv^T)x &= x - 2\frac{(x + \|x\|y)(x^T x + \|x\|x^T y)}{\|x\|^2 + 2x^T y\|x\| + \|x\|^2\|y\|^2} \\ &= x - (x - \|x\|y) \\ &= \|x\|y.\end{aligned}$$

If we use $y = \pm e_1$, we can get a reflection that zeros out all but the first element of the vector x . So with appropriate choices of reflections, we can take a matrix A and zero out all of the subdiagonal elements of the first column.




Figure 1: Construction of a reflector to transform x into $\|x\|y$ where $y = 1$.

Now think about applying a sequence of Householder transformations to introduce subdiagonal zeros into A , just as we used a sequence of Gauss transformations to introduce subdiagonal zeros in Gaussian elimination. As with LU factorization, we can re-use the storage of A by recognizing that the number of nontrivial parameters in the vector w at each step is the same as the number of zeros produced by that transformation. This gives us the following:

```
function hqr!(A)
    m,n = size(A)
    τ = zeros(n)

    for j = 1:n

        # Find H = I-τ*w*w' to zero out A[j+1:end,j]
        normx = norm(A[j:end,j])
        s = -sign(A[j,j])
        u1 = A[j,j] - s*normx
        w = A[j:end,j]/u1
        w[1] = 1.0
        A[j+1:end,j] = w[2:end] # Save trailing part of w
        A[j,j] = s*normx # Diagonal element of R
```

```

    tau[j] = -s*u1/normx      # Save scaling factor

    # Update trailing submatrix by multiplying by H
    A[j:end,j+1:end] -= tau[j]*w*(w'*A[j:end,j+1:end])

end

A, tau
end

```

If we ever need Q or Q^T explicitly, we can always form it from the compressed representation. We can also multiply by Q and Q^T implicitly:

```

function applyQ!(QR, τ, X)
    m, n = size(QR)
    for j = n:-1:1
        w = [1.0; QR[j+1:end,j]]
        X[j:end,:] -= tau[j]*w*(w'*X[j:end,:])
    end
    X
end

function applyQT!(QR, τ, X)
    m, n = size(QR)
    for j = 1:n
        w = [1.0; QR[j+1:end,j]]
        X[j:end,:] -= tau[j]*w*(w'*X[j:end,:])
    end
    X
end

applyQ(QR, tau, X) = applyQ!(QR, tau, copy(X))
applyQT(QR, tau, X) = applyQ(QR, tau, copy(X))

```

Givens rotations

Householder reflections are one of the standard orthogonal transformations used in numerical linear algebra. The other standard orthogonal transformation is a *Givens rotation*:

$$G = \begin{bmatrix} c & -s \\ s & c \end{bmatrix}.$$

where $c^2 + s^2 = 1$. Note that

$$G = \begin{bmatrix} c & -s \\ s & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} cx - sy \\ sx + cy \end{bmatrix}$$

so if we choose

$$s = \frac{-y}{\sqrt{x^2 + y^2}}, \quad c = \frac{x}{\sqrt{x^2 + y^2}}$$

then the Givens rotation introduces a zero in the second column. More generally, we can transform a vector in \mathbb{R}^m into a vector parallel to e_1 by a sequence of $m - 1$ Givens rotations, where the first rotation moves the last element to zero, the second rotation moves the second-to-last element to zero, and so forth.

For some applications, introducing zeros one by one is very attractive. In some places, you may see this phrased as a contrast between algorithms based on Householder reflections and those based on Givens rotations, but this is not quite right. Small Householder reflections can be used to introduce one zero at a time, too. Still, in the general usage, Givens rotations seem to be the more popular choice for this sort of local introduction of zeros.