

Week 6: Wednesday, Sep 28

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.

Now think about applying a sequence of Householder transformations to introduce subdiagonal zeros into A , just as we used a sequence of Gauss




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

transformations to introduce subdiagonal zeros in Gaussian elimination. This leads us to the following algorithm to compute the QR decomposition:

```
function [Q,R] = lec16hqr1(A)
% Compute the QR decomposition of an m-by-n matrix A using
% Householder transformations.

[m,n] = size(A);
Q = eye(m);      % Orthogonal transform so far
R = A;          % Transformed matrix so far

for j = 1:n

    % -- Find H = I-tau*w*w' to put zeros below R(j,j)
    normx = norm(R(j:end,j));
    s      = -sign(R(j,j));
    u1     = R(j,j) - s*normx;
    w      = R(j:end,j)/u1;
    w(1)   = 1;
    tau    = -s*u1/normx;

    % -- R := HR, Q := QH
    R(j:end,:) = R(j:end,:)-(tau*w)*(w'*R(j:end,:));
    Q(:,j:end) = Q(:,j:end)-(Q(:,j:end)*w)*(tau*w)';
end
```

end

Note that there are two valid choices of u_1 at each step; we make the choice that avoids cancellation in the obvious version of the formula.

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 [A,tau] = lec16hqr2(A)
% Compute the QR decomposition of an m-by-n matrix A using
% Householder transformations, re-using the storage of A
% for the Q and R factors.

[m,n] = size(A);
tau = zeros(n,1);

for j = 1:n

    % -- Find H = I-tau*w*w' to put zeros below A(j,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;
    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;

    % -- R := HR
    A(j:end,j+1:end) = A(j:end,j+1:end)-...
        (tau(j)*w)*(w'*A(j:end,j+1:end));

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 QX = lec16applyQ(QR,tau,X)

    [m,n] = size(QR);
    QX = X;
    for j = n:-1:1
        w = [1; QR(j+1:end,j)];
        QX(j:end,:) = QX(j:end,)-(tau(j)*w)*(w'*QX(j:end,:));
    end
```

```
function QTX = lec16applyQT(QR,tau,X)

    [m,n] = size(QR);
    QTX = X;
    for j = 1:n
        w = [1; QR(j+1:end,j)];
        QTX(j:end,:) = QTX(j:end,)-(tau(j)*w)*(w'*QTX(j:end,:));
    end
```

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.