Week 11: Wednesday, Nov 4

Logistics

1. HW 4 is due in one week.

2. I will be traveling from the afternoon of Nov 11 through the morning of Nov 14. Charlie Van Loan will teach next Friday's lecture.

Recap

Our basic QR algorithm had two serious flaws: each iteration was slow, and convergence could take many iterations. Last time, we showed how converting to Hessenberg form made each step of the QR iteration run in $O(n^2)$ time. We now turn to the problem of convergence.

Inverse iteration and the QR method

When we discussed the power method, we found that we could improve convergence by a spectral transformation that mapped the eigenvalue we wanted to something with large magnitude (preferably much larger than the other eigenvalues). This was the *shift-invert* strategy. We already know there is a connection leading from the power method to orthogonal iteration to the QR method, which we can summarize with a small number of formulas. Let us see if we can follow the same path to uncover a connection from inverse iteration (the power method with A^{-1}, a special case of shift-invert in which the shift is zero) to QR. If we call the orthogonal factors in orthogonal iteration $\underline{Q}^{(k)}$ ($\underline{Q}^{(0)} = I$) and the iterates in QR iteration $A^{(k)}$, we have

$$(1) \qquad\qquad A^k = \underline{Q}^{(k)}\underline{R}^{(k)}$$
$$(2) \qquad\qquad A^{(k)} = (\underline{Q}^{(k)})^* A(\underline{Q}^{(k)}).$$

In particular, note that because $R^{(k)}$ are upper triangular,

$$A^k e_1 = (\underline{Q}^{(k)} e_1) r_{11}^{(k)};$$

that is, the first column of $\underline{Q}^{(k)}$ corresponds to the kth step of power iteration starting at e_1. What happens when we consider negative powers of A? Inverting (1), we find

$$A^{-k} = (\underline{R}^{(k)})^{-1}(\underline{Q}^{(k)})^*$$

The matrix $\tilde{R}^{(k)} = (\underline{R}^{(k)})^{-1}$ is again upper triangular; and if we look carefully, we can see in this fact another power iteration:

$$e_n^* A^{-k} = e_n^* \tilde{R}^{(k)}(\underline{Q}^{(k)})^* = \tilde{r}_{nn}^{(k)}(\underline{Q}^{(k)}e_n)^*.$$

That is, the last column of $\underline{Q}^{(k)}$ corresponds to a power iteration converging to a *row* eigenvector of A^{-1}.

Shifting gears

The connection from inverse iteration to orthogonal iteration (and thus to QR iteration) gives us a way to incorporate the shift-invert strategy into QR iteration: simply run QR on the matrix $A - \sigma I$, and the (n, n) entry of the iterates (which corresponds to a Rayleigh quotient with an increasingly-good approximate row eigenvector) should start to converge to $\lambda - \sigma$, where λ is the eigenvalue nearest σ. Put differently, we can run the iteration:

$$Q^{(k)}R^{(k)} = A^{(k-1)} - \sigma I$$
$$A^{(k)} = R^{(k)}Q^{(k)} + \sigma I.$$

If we choose a good shift, then the lower right corner entry of $A^{(k)}$ should converge to the eigenvalue closest to σ in fairly short order, and the rest of the elements in the last row should converge to zero.

The shift-invert power iteration converges fastest when we choose a shift that is close to the eigenvalue that we want. We can do even better if we choose a shift *adaptively*, which was the basis for running Rayleigh quotient iteration. The same idea is the basis for the *shifted QR iteration*:

(3) $$Q^{(k)}R^{(k)} = A^{(k-1)} - \sigma_k I$$

(4) $$A^{(k)} = R^{(k)}Q^{(k)} + \sigma_k I.$$

This iteration is equivalent to computing

$$\underline{Q}^{(k)}\underline{R}^{(k)} = \prod_{j=1}^{n}(A - \sigma_j I)$$
$$A^{(k)} = (\underline{Q}^{(k)})^* A (\underline{Q}^{(k)})$$
$$\underline{Q}^{(k)} = Q^{(k)} Q^{(k-1)} \cdots Q^{(1)}.$$

What should we use for the shift parameters σ_k? A natural choice is to use $\sigma_k = e_n^* A^{(k-1)} e_n$, which is the same as $\sigma_k = (\underline{Q}^{(k)} e_n)^* A (\underline{Q}^{(k)} e_n)$, the Rayleigh quotient based on the last column of $\underline{Q}^{(k)}$. This simple shifted QR iteration is equivalent to running Rayleigh iteration starting from an initial vector of e_n, which we noted before is locally quadratically convergent.

Double trouble

The simple shift strategy we described in the previous section gives *local* quadratic convergence, but it is not *globally* convergent. As a particularly pesky example, consider what happens if we want to compute a complex conjugate pair of eigenvalues of a real matrix. With our simple shifting strategy, the iteration (3)–(4) will never produce a complex iterate, a complex shift, or a complex eigenvalue. The best we can hope for is that our initial shift is closer to both eigenvalues in the conjugate pair than it is to anything else in the spectrum; in this case, we will most likely find that the last two columns of $\underline{Q}^{(k)}$ are converging to a basis for an *invariant row subspace* of A, and the corresponding eigenvalues are the eigenvalues of the trailing 2-by-2 sub-block.

Fortunately, we know how to compute the eigenvalues of a 2-by-2 matrix! This suggests the following shift strategy: let σ_k be one of the eigenvalues of $A^{(k)}(n-1:n, n-1:n)$. Because this 2-by-2 problem can have complex roots even when the matrix is real, this shift strategy allows the possibility that we could converge to complex eigenvalues. On the other hand, if our original matrix is real, perhaps we would like to consider the *real* Schur form, in which U is a real matrix and T is block diagonal with 1-by-1 and 2-by-2 diagonal blocks that correspond, respectively, to real and complex eigenvalues. If we

shift with *both* roots of $A^{(k)}(n-1:n, n-1:n)$, equivalent to computing

$$Q^{(k)}R^{(k)} = (A^{(k-1)} - \sigma_{k+}I)(A^{(k-1)} - \sigma_{k-})$$
$$A^{(k)} = (Q^{(k)})^* A^{(k-1)} Q^{(k)}.$$

There is one catch here: even if we started with $A^{(0)}$ in Hessenberg form, it is unclear how to do this double-shift step in $O(n^2)$ time!

The following fact will prove our salvation: if we Q and V are both orthogonal matrices and $Q^T A Q$ and $V^T A V$ are both (unreduced) Hessenberg[1]) and the first column of Q is the same as the first column of V, then all successive columns of Q are unit scalar multiples of the corresponding columns of V. This is the *implicit Q theorem*. Practically, it means that we can do any sort of shifted QR step we would like in the following way:

1. Apply as a similarity any transformations in the QR decomposition that affect the leading submatrix (1-by-1 or 2-by-2).

2. Restore the resulting matrix to Hessenberg form without further transformations to the leading submatrix.

In the first step, we effectively compute the first column of Q; in the second step, we effectively compute the remaining columns. Certainly we compute *some* transformation with the right leading column; and the implicit Q theorem tells us that any such transformation is basically the one we would have computed with an ordinary QR step.

---

[1] An unreduced Hessenberg matrix has no zeros on the first subdiagonal.