

Convergence of Eigenvectors

A fundamental question posed by power (or inverse power) iterations is “when do we stop?” A naive criterion might be to stop when further iteration provides little change. That is fine, but we still need to define “little”, and we would also like to have some indication of how accurate our “converged” eigenpair is. So let’s define the residual of an approximate eigenpair (λ, v) to be

$$r = r(A, \lambda, v) = Av - \lambda v.$$

Now if (λ, v) is an exact eigenpair for A , then $r = 0$, so what about small r ? Let $\|v\|_2 = 1$ and define $\epsilon = \|r\|_2$, and consider the matrix $A - rv^t$:

$$(A - rv^t)v = Av - r = Av - (Av - \lambda v) = \lambda v.$$

Thus (λ, v) is an exact eigenpair for the matrix $A - rv^t$, and (λ, v) must be a backward stable eigenpair if, e.g. $\|rv^t\|_2/\|A\|_2 = \epsilon/\|A\|_2$ is small. In an iteration for which r can be computed without excessive cost, this gives a rational stopping criterion.

If we are doing an inverse iteration, then x_{j+1} is defined by $(A - sI)x_{j+1} = x_j$, and the approximate eigenvalue, say l , of $(A - sI)^{-1}$, is related to λ as $\lambda = s + 1/l$. We want a residual for A (not for $(A - sI)^{-1}$ (right?)), so our residual is

$$\begin{aligned} r &= Ax_{j+1} - \lambda x_{j+1} \\ &= x_j + sx_{j+1} - \lambda x_{j+1} \\ &= x_j - (1/l)x_{j+1}. \end{aligned}$$

However we come to the residual $r(A, \lambda, v)$, we can bound the error in the computed eigenvalue λ . The relative condition number for λ is $\kappa(\lambda) = \frac{\|A\|}{|\lambda| \|w^*v\|}$, where v is as above, and w is the corresponding normalized left eigenvector:

$w^*A = \bar{w}^t A = \lambda w^*$, $\|w\|_2 = 1$. Thus we can expect that A has an eigenvalue μ satisfying

$$|\mu - \lambda| \leq |\mu| \kappa(\mu) \frac{\|r\|_2}{\|A\|_2} = \frac{|\epsilon|}{|w^*v|}.$$

In practice, we can compute w with one (left) inverse iteration with shift $s = \lambda$.

Worried that inverse iteration with an accurate shift is ineffective because of illconditioning in $A - sI$? Good, you should be. It turns out, though, that the worry is misplaced. Here is the outline: If s is a good approximation to an eigenvalue of A , then $1/(\lambda - s)$ is a good approximation to the largest eigenvalue of $(A - sI)^{-1}$.

Recall the parameterized system modeling the errors in solving $(A - sI)x = b$:

$$(A - sI + tE)x(t) = b + te.$$

The (first order in t) error in x is

$$\dot{x}(0) = (A - sI)^{-1}(e - E(A - sI)^{-1}b) \equiv (A - sI)^{-1}c.$$

The quantity we want is $(A - sI)^{-1}b$ and our error is $(A - sI)^{-1}c + O(t^2)$, giving a *computed* x with its error almost entirely in direction we are trying to find! The better s , the more illconditioned $A - sI$, but the better is our inverse iteration!