Vol. 22, No. 5, October 1985

006

RESIDUAL INVERSE ITERATION FOR THE NONLINEAR EIGENVALUE PROBLEM*

A. NEUMAIER†

Abstract. For the nonlinear eigenvalue problem $A(\hat\lambda)\hat x = 0$, where $A(\cdot)$ is a matrix-valued operator, residual inverse iteration with shift σ is defined by

$$x^{(l+1)} := \text{const.}\,(x^{(l)} - A(\sigma)^{-1}A(\lambda_{l+1})x^{(l)}),$$

where λ_{l+1} is an appropriate approximation of $\hat\lambda$. In the linear case, $A(\lambda) = A - \lambda I$, this is theoretically equivalent to ordinary inverse iteration, but the residual formulation results in a considerably higher limit accuracy when the residual $A(\lambda_{l+1})x^{(l)} = Ax^{(l)} - \lambda_{l+1}x^{(l)}$ is accumulated in double precision. In the nonlinear case, if σ is sufficiently close to $\hat\lambda$, convergence is at least linear with convergence factor proportional to $|\sigma - \hat\lambda|$. As with ordinary inverse iteration, the convergence can be accelerated by using variable shifts.

1. Introduction. Inverse iteration is generally considered as one of the standard methods for the computation of selected eigenpairs of a linear eigenvalue problem. Although it is best analyzed in terms of eigenvector expansions (see e.g. Wilkinson [11]), local convergence can also be proved by deriving inverse iteration from Newton's method applied to a suitable equivalent system of nonlinear equations (Unger [10]). To treat the more general nonlinear eigenvalue problem, this latter approach can be generalized and leads to a nonlinear version of inverse iteration for the eigenvalue problem $A(\lambda)x = 0$, $x \neq 0$, namely

$$\begin{aligned} y^{(l)} &= A(\lambda_l)^{-1}A'(\lambda_l)x^{(l)}, \\ x^{(l+1)} &= y^{(l)}/e^*y^{(l)}, \qquad \lambda_{l+1} = \lambda_l - 1/e^*y^{(l)}. \end{aligned} \tag{1}$$

The numerical behaviour of this and other related methods is discussed in the survey by Ruhe [6]. An essential disadvantage of these methods is the fact that in each step the coefficient matrix $A(\lambda_l)$ of the linear system for $y^{(l)}$ changes; and, in contrast to the linear case, working with a fixed "shift" σ instead of λ_l results in convergence to the wrong problem, namely to a solution of the linearized problem $A(\lambda^*)x = (\lambda^* - \sigma)A'(\lambda^*)x$.

In the following, we circumvent this difficulty by considering a variant of inverse iteration based on the use of the residual. To motivate the new approach we rearrange (1) such that

$$x^{(l+1)} = x^{(l)} - dx^{(l)}$$

is computed from $x^{(l)}$ by subtracting the correction term

$$\begin{aligned} dx^{(l)} = x^{(l)} - x^{(l+1)} &= x^{(l)} + (\lambda_{l+1} - \lambda_l)y^{(l)} \\ &= x^{(l)} + (\lambda_{l+1} - \lambda_l)A(\lambda_l)^{-1}A'(\lambda_l)x^{(l)} \\ &= A(\lambda_l)^{-1}(A(\lambda_l) + (\lambda_{l+1} - \lambda_l)A'(\lambda_l))x^{(l)} \\ &= A(\lambda_l)^{-1}A(\lambda_{l+1})x^{(l)} + O((\lambda_{l+1} - \lambda_l)^2), \end{aligned}$$

if $A(\lambda)$ is twice continuously differentiable. By neglecting the error term we obtain for $x^{(l+1)}$ the expression

$$x^{(l+1)} = x^{(l)} - A(\lambda_l)^{-1}A(\lambda_{l+1})x^{(l)}, \tag{2}$$

* Received by the editors February 21, 1984, and in revised form September 7, 1984.
† Institut für Angewandte Mathematik, Universität Freiburg, D-7800 Freiburg, West Germany.

where the new approximation λ_{l+1} for the eigenvalue now has to be determined beforehand (we shall use a generalized Rayleigh quotient). It turns out that in this formulation λ_l may be replaced by a constant "shift" σ without destroying convergence to the wanted eigenpair.

The new algorithm, called residual inverse iteration, thus computes the new approximation $x^{(l+1)}$ by applying to $x^{(l+1)}$ a correction term computed from the residual $A(\lambda)x^{(l)}$ for a suitable λ. Hence in the presence of rounding errors residual inverse iteration with double precision accumulation of the residuals gives about the same limit accuracy as one would get with ordinary inverse iteration only when the complete iteration is performed in double precision. In particular, for linear problems, residual inverse iteration can be profitably used to refine eigenvalue approximations obtained from the QR or QZ algorithm (see e.g. Stewart [8]) since the single precision partial factorization available from the QR or QZ algorithm can be reused to save factorization time. Thus residual inverse iteration provides a simple alternative to some refinement procedures proposed in the literature ([1], [2, § 62], [9], [12]), and has the advantage of preserving the structure of A and not requiring an initial eigenvector approximation.

In case that $A(\lambda) = A - \lambda I$, residual inverse iteration is again theoretically equivalent to ordinary inverse iteration. But in the nonlinear case, residual inverse iteration is no longer strictly equalent to (1), and can be used either with a fixed shift σ or with variable shift. For the fixed shift, local convergence is at least linear with a convergence factor proportional to the distance of σ to the nearest eigenvalue $\hat\lambda$ (provided that $\hat\lambda$ is simple and isolated). Double precision computation of the residuals again leads (in well-conditioned cases) to results which are correct to almost double precision.

The paper is organized as follows. In § 2, residual inverse iteration is defined for fixed shift. Section 3 gives the local convergence proof, with some remarks on the convergence behaviour in case of variable shifts. In § 4, we comment on the practical realization and demonstrate the behaviour of the algorithm with three examples: the Frank matrix of order 11 and two definite quadratic eigenvalue problems of Scott and Ward [7]. We use the notation $\mathbb{C}^{n\times n}$ for the set of complex square $n\times n$-matrices, denote conjugate transposition by an asterisk *, and use $\|\cdot\|$ for an arbitrary vector norm.

2. The algorithm. We consider the finite-dimensional nonlinear eigenvalue problem

$$A(\hat\lambda)\hat x = 0, \quad \hat\lambda \in D \subseteq \mathbb{C}, \quad \hat x \in \mathbb{C}^n - \{0\}, \tag{3}$$

where $A: D \to \mathbb{C}^{n\times n}$ is a continuous matrix-valued map. We suppose that an approximation $\sigma \in D$ to $\hat\lambda$ is known, that $A(\sigma)$ is nonsingular, and that e is a normalization vector such that

$$e^*\hat x = 1; \tag{4}$$

usually e will be the unit vector with a 1 in the position of the largest entry of $\hat x$. We suggest the following iteration for the approximation of a solution of (3).

Residual inverse iteration.

Step 1. Put $l = 0$, and compute an initial approximation $x^{(0)}$ to $\hat x$ as the normalized solution of the equation

$$A(\sigma)\bar x^{(0)} = b, \qquad x^{(0)} := \bar x^{(0)}/e^*\bar x^{(0)}; \tag{5}$$

the vector $b \neq 0$ has to be chosen suitably (see § 4).

Step 2. Compute an improved approximation λ_{l+1} to $\hat{\lambda}$ by solving one of the equations

$$x^{(l)*}A(\lambda_{l+1})x^{(l)}=0 \quad \text{or} \tag{6a}$$

$$e^*A(\sigma)^{-1}A(\lambda_{l+1})x^{(l)}=0. \tag{6b}$$

Formula (6a) is appropriate only when $A(\lambda)$ is Hermitian and λ_l is real; otherwise, (6b) has to be used. The root closest to λ_l is accepted as λ_{l+1}.

Step 3. Compute the residual

$$r^{(l)}:=A(\lambda_{l+1})x^{(l)}. \tag{7}$$

Step 4. Compute an improved approximation $x^{(l+1)}$ to $\hat{x}$ by solving the equation

$$A(\sigma)\,dx^{(l)}=r^{(l)} \tag{8}$$

and normalizing the vector

$$\bar{x}^{(l+1)}:=x^{(l)}-dx^{(l)}, \qquad x^{(l+1)}:=\bar{x}^{(l+1)}/e^*\bar{x}^{(l+1)}. \tag{9}$$

Step 5. Increase l by one and return to Step 2.

In the special case $A(\lambda)=A-\lambda I$, residual inverse iteration is equivalent to ordinary inverse iteration with shift σ (in the absence of rounding errors); indeed we then have

$$\begin{aligned}(A-\sigma I)\bar{x}^{(l+1)}&=(A-\sigma I)x^{(l)}-(A-\sigma I)\,dx^{(l)}\\&=(A-\sigma I)x^{(l)}-(A-\lambda_{l+1}I)x^{(l)}\\&=(\lambda_{l+1}-\sigma)x^{(l)},\end{aligned}$$

so that $\bar{x}^{(l+1)}$, and hence $x^{(l+1)}$, is parallel to $(A-\sigma I)^{-1}x^{(l)}$. Thus we can hope that in the more general situation discussed above, some of the excellent convergence properties of inverse iteration (as discussed e.g. in Wilkinson [11], Parlett [5]) are still valid.

3. Convergence analysis. In this section we shall assume that the matrix function $A(\lambda)$ is twice continuously differentiable in some neighbourhood U of $\hat{\lambda}$. Then the *divided difference*

$$A[\lambda_1,\lambda_2]:=\begin{cases}\dfrac{A(\lambda_2)-A(\lambda_1)}{\lambda_2-\lambda_1} & \text{if } \lambda_1\neq\lambda_2,\\ A'(\lambda_1) & \text{if } \lambda_1=\lambda_2,\end{cases}$$

is defined in U, is continuously differentiable and satisfies the relations

$$A(\lambda_2)=A(\lambda_1)+(\lambda_2-\lambda_1)A[\lambda_1,\lambda_2], \qquad A[\lambda_1,\lambda_2]=A'(\hat{\lambda})+O(\lambda_1-\hat{\lambda},\lambda_2-\hat{\lambda}).$$

If $A(\hat{\lambda})$ is singular we call each vector $\hat{x}\neq 0$ satisfying $A(\hat{\lambda})\hat{x}=0$ a *right eigenvector* associated with $\hat{\lambda}$, and each vector $\hat{y}\neq 0$ satisfying $\hat{y}^*A(\hat{\lambda})=0$ a *left eigenvector* associated with $\hat{\lambda}$.

PROPOSITION 1. *The following conditions are equivalent:*

(i) $d(\lambda):=\det A(\lambda)$ *has a simple zero at* $\lambda=\hat{\lambda}$;

(ii) $A(\hat{\lambda})$ *has corank* 1, *and for any pair* $\hat{x}$, $\hat{y}$ *of right and left eigenvectors,*

$$\hat{y}^*A'(\hat{\lambda})\hat{x}\neq 0. \tag{10}$$

Proof. Suppose first that $A(\hat{\lambda})$ has corank 1. Since the adjoint matrix $C:=\operatorname{Adj}A(\hat{\lambda})$ satisfies $CA(\hat{\lambda})=A(\hat{\lambda})C=\det(A(\hat{\lambda}))I=0$, the columns of C are multiples of $\hat{x}$, and the rows of C are multiples of $\hat{y}^*$. Therefore $C=\gamma\hat{x}\hat{y}^*$ for a suitable constant γ, and since some $(n-1)\times(n-1)$ minor of $A(\hat{\lambda})$ is nonzero, $C\neq 0$ whence $\gamma\neq 0$. Now by Gröbner [3, eq. (4.76)],

$$\begin{aligned}d(\lambda)=\det A(\lambda)&=\det(A(\hat{\lambda})+(\lambda-\hat{\lambda})A[\lambda,\hat{\lambda}])\\&=\det A(\hat{\lambda})+(\lambda-\hat{\lambda})\operatorname{tr}(\operatorname{Adj}A(\hat{\lambda})\cdot A[\lambda,\hat{\lambda}])+O(\lambda-\hat{\lambda})^2\\&=(\lambda-\hat{\lambda})\operatorname{tr}(CA[\lambda,\hat{\lambda}])+O(\lambda-\hat{\lambda})^2\\&=(\lambda-\hat{\lambda})\operatorname{tr}(\gamma\hat{x}\hat{y}^*A'(\hat{\lambda}))+O(\lambda-\hat{\lambda})^2\\&=(\lambda-\hat{\lambda})\gamma\cdot\hat{y}^*A'(\hat{\lambda})\hat{x}+O(\lambda-\hat{\lambda})^2.\end{aligned}$$

Hence in this case, (i) and (ii) are equivalent.

Suppose now that $A(\hat{\lambda})$ has corank $s\neq 1$. If $s=0$ then $d(\hat{\lambda})\neq 0$ and neither (i) nor (ii) holds. And if $s\geqq 2$ then all $(n-1)\times(n-1)$ minors of $A(\hat{\lambda})$ are zero whence $C=0$ and, as above, $d(\lambda)=O(\lambda-\hat{\lambda})^2$. Again neither (i) nor (ii) holds. This proves the proposition. □

We shall call $\hat{\lambda}$ a *simple isolated eigenvalue* of the matrix function $A(\lambda)$ if $A(\lambda)$ is twice continuously differentiable in some neighbourhood of $\hat{\lambda}$ and the conditions (i) and (ii) of Proposition 1 are satisfied.

PROPOSITION 2. *Let* $\hat{\lambda}$ *be a simple isolated eigenvalue of* $A(\lambda)$, *and let* $\hat{x}$ *be a corresponding right eigenvector normalized such that* $e^*\hat{x}=1$. *Then the matrix*

$$B:=A(\hat{\lambda})+A'(\hat{\lambda})\hat{x}e^* \tag{11}$$

is nonsingular.

Proof. Assume that $Bx=0$. Then, with a left eigenvector $\hat{y}$, we have $0=\hat{y}^*Bx=\hat{y}^*A(\hat{\lambda})x+\hat{y}^*A'(\hat{\lambda})\hat{x}e^*x=\hat{y}^*A'(\hat{\lambda})\hat{x}\cdot e^*x$, and by (10) then $e^*x=0$. Therefore $A(\hat{\lambda})x=Bx-A'(\hat{\lambda})\hat{x}e^*x=0$, and $x=t\hat{x}$ for suitable t since $A(\hat{\lambda})$ has corank 1. Now $t=te^*\hat{x}=e^*x=0$ implies $x=0$. Since x was arbitrary, B is nonsingular. □

PROPOSITION 3. *With the assumptions of Proposition* 2, *suppose that for sufficiently small* $\delta\geqq\varepsilon>0$ *we have* $0<|\sigma-\hat{\lambda}|\leqq\delta$ *and*

$$\lambda=\hat{\lambda}+O(\varepsilon), \qquad x=\hat{x}+O(\varepsilon).$$

Then $A(\sigma)$ *is nonsingular, and if* $\sigma\neq\lambda$ *then the vector*

$$\bar{x}:=x-A(\sigma)^{-1}A(\lambda)x \tag{12}$$

satisfies

$$0\neq e^*\bar{x}=\frac{\sigma-\lambda}{\sigma-\hat{\lambda}}(1+O(\varepsilon)), \tag{13}$$

$$\bar{x}/e^*\bar{x}=\hat{x}+(\sigma-\hat{\lambda})O(\varepsilon). \tag{14}$$

Proof. If $\sigma\neq\hat{\lambda}$ is in a sufficiently small neighbourhood of the simple eigenvalue $\hat{\lambda}$ then $\det A(\sigma)\neq 0$ whence $A(\sigma)$ is nonsingular. Define

$$S:=A(\sigma)+(1-\sigma+\hat{\lambda})A[\sigma,\hat{\lambda}]\hat{x}e^*. \tag{15}$$

Then, with B defined by (13),

$$\begin{aligned}S&=B+A(\sigma)-A(\hat{\lambda})+(A[\sigma,\lambda]-A'(\hat{\lambda})-(\sigma-\hat{\lambda})A[\sigma,\hat{\lambda}])\hat{x}e^*\\&=B+O(\sigma-\hat{\lambda}),\end{aligned}$$

and since B is nonsingular by Proposition 2, S is nonsingular, and

$$S^{-1} = B^{-1} + O(\sigma - \hat{\lambda}) = O(1). \tag{16}$$

Moreover,

$$\begin{aligned} S\hat{x} &= (A(\sigma) - A(\hat{\lambda}))\hat{x} + (1 - \sigma + \hat{\lambda})A[\sigma, \hat{\lambda}]\hat{x}e^*\hat{x} \\ &= (\sigma - \hat{\lambda})A[\sigma, \hat{\lambda}]\hat{x} + (1 - \sigma + \hat{\lambda})A[\sigma, \hat{\lambda}]\hat{x}, \end{aligned}$$

so that

$$A[\sigma, \hat{\lambda}]\hat{x} = S\hat{x}. \tag{17}$$

Since $e^*(\bar{x} - e^*\bar{x} \cdot \hat{x}) = e^*\bar{x} - e^*\bar{x} \cdot e^*\hat{x} = 0$, (15), (12), and (17) imply

$$\begin{aligned} z &:= S(\bar{x} - e^*\bar{x} \cdot \hat{x}) = A(\sigma)(x - e^*\bar{x} \cdot \hat{x}) \\ &= A(\sigma)\bar{x} - e^*\bar{x} \cdot A(\sigma)\hat{x} \\ &= (A(\sigma) - A(\lambda))x - e^*\bar{x}(A(\sigma) - A(\hat{\lambda}))\hat{x} \\ &= (\sigma - \lambda)A[\sigma, \lambda]x - e^*\bar{x}(\sigma - \hat{\lambda})A[\sigma, \hat{\lambda}]\hat{x} \\ &= (\sigma - \lambda)(A[\sigma, \hat{\lambda}]\hat{x} + O(\varepsilon)) - e^*\bar{x}(\sigma - \hat{\lambda})A[\sigma, \hat{\lambda}]\hat{x} \\ &= (\sigma - \lambda)(S\hat{x} + O(\varepsilon)) - e^*\bar{x}(\sigma - \hat{\lambda})S\hat{x}. \end{aligned}$$

By (16), this implies

$$\bar{x} - e^*\bar{x} \cdot \hat{x} = S^{-1}z = (\sigma - \lambda)(\hat{x} + O(\varepsilon)) - e^*\bar{x}(\sigma - \hat{\lambda})\hat{x}. \tag{18}$$

Multiplication with e^* gives

$$0 = (\sigma - \lambda)(1 + O(\varepsilon)) - e^*\bar{x}(\sigma - \hat{\lambda})$$

which implies (13) and $\sigma - \lambda = e^*\bar{x}(\sigma - \hat{\lambda})(1 + O(\varepsilon))$. Insertion into (18) and division by $e^*\bar{x}$ finally gives $\bar{x}/e^*\bar{x} - \hat{x} = (\sigma - \hat{\lambda})O(\varepsilon)$ which implies (14). □

PROPOSITION 4. *Under the assumptions of Proposition 2, if, for sufficiently small $\delta \geq \varepsilon > 0$ we have $0 < |\sigma - \hat{\lambda}| \leq \delta$ and $x^{(l)} = \hat{x} + O(\varepsilon)$, then the zero of (6b) closest to $\hat{\lambda}$ satisfies*

$$\lambda_{l+1} = \hat{\lambda} + O(\varepsilon); \tag{19}$$

and in case that $A(\lambda)$ is Hermitian and $\hat{\lambda}$ is real then the zero of (6a) closest to $\hat{\lambda}$ satisfies

$$\lambda_{l+1} = \hat{\lambda} + O(\varepsilon^2). \tag{20}$$

Proof. Write $\bar{y} = e^*A(\sigma)^{-1}$, $y = \bar{y}/\|\bar{y}\|$, and consider the function $f(\lambda) := y^*A(\lambda)x^{(l)}$. For $\delta, \varepsilon \to 0$, $A(\sigma)$ approaches $A(\hat{\lambda})$ whence y approaches a left nullvector $\hat{y}$ of $A(\hat{\lambda})$, i.e. a left eigenvector corresponding to $\hat{\lambda}$, and $f(\lambda)$ approaches the function $\hat{f}(\lambda) := \hat{y}^*A(\lambda)\hat{x}$ which has a simple zero at $\hat{\lambda}$. Therefore if δ is sufficiently small, $f(\lambda)$ has a simple zero λ_{l+1} close to $\hat{\lambda}$. Now $0 = f(\lambda_{l+1}) = f(\hat{\lambda}) + (\lambda_{l+1} - \hat{\lambda})f'(\xi)$ for some ξ near $\hat{\lambda}$ whence

$$\lambda_{l+1} = \hat{\lambda} + \frac{f(\hat{\lambda})}{f'(\xi)} = \hat{\lambda} + O(f(\hat{\lambda}))$$

since $f'(\xi)$ is bounded away from zero (near $\hat{\lambda}$). But

$$f(\hat{\lambda}) = y^*A(\hat{\lambda})x^{(l)} = y^*A(\hat{\lambda})(x^{(l)} - \hat{x}) = O(\varepsilon),$$

so that (19) holds. (20) is proved in the same way with $\bar{y} = x^{(l)}$, observing that in the Hermitian case $\hat{y} := \hat{x}$ is a left eigenvector, and

$$f(\hat{\lambda}) = x^{(l)*}A(\hat{\lambda})x^{(l)} = (x^{(l)} - \hat{x})^*A(\hat{\lambda})(x^{(l)} - \hat{x}) = O(\varepsilon^2). \qquad \square$$

With the observation that for $\sigma \to \hat{\lambda}$ the solution $x^{(0)}$ of (5) converges to $\hat{x}$, Propositions 3 and 4 lead to the following local convergence theorem.

THEOREM. *Let $\hat{\lambda}$ be a simple isolated eigenvalue of $A(\lambda)$, and suppose that there is a corresponding eigenvector $\hat{x}$ normalized such that $e^*\hat{x} = 1$. Then the residual inverse iteration converges for all σ sufficiently close to $\hat{\lambda}$, and we have*

$$\frac{\|x^{(l+1)} - \hat{x}\|}{\|x^{(l)} - \hat{x}\|} = O(\sigma - \hat{\lambda}), \qquad \|\lambda^{(l+1)} - \hat{\lambda}\| = O(\|x^{(l)} - \hat{x}\|^t),$$

where $t = 1$ if (6b) is used, and $t = 2$ if $A(\lambda)$ is Hermitian, $\hat{\lambda}$ is real, and (6a) is used. □

The theorem implies local linear convergence with a convergence factor proportional to $|\sigma - \hat{\lambda}|$. In particular, this suggests that the convergence is accelerated by updating the shift σ in each iteration step (or in some iteration steps only, if the extra work to refactor $A(\sigma)$ is considered as being too much). It follows easily from the theorem that we have quadratic convergence (and in the Hermitian case with real $\hat{\lambda}$ even cubic convergence) if in each iteration step, σ is replaced by the most recent value of λ_l.

4. Numerical examples. For the actual computation on a computer, several remarks are in place. (5) is usually solved by using a factorization

$$A(\sigma) = SR, \tag{21a}$$

where R is upper triangular, and S is a permuted lower unit triangular or orthogonal matrix. An appropriate choice of b is then the vector $b = Sj$, $j = (1, \ldots, 1)^*$, so that we actually solve

$$R\bar{x}^{(0)} = (1, \cdots, 1)^*, \qquad x^{(0)} = \bar{x}^{(0)}/e^*\bar{x}^{(0)} \tag{5a}$$

in place of (5). This choice is motivated in Wilkinson [11] for ordinary inverse iteration and works well in the present algorithm.

In the special case that $A(\lambda)$ is linear in λ, and the QR (or QZ) algorithm has been used to compute the eigenvalues (cf. Wilkinson [11], Parlett [5], Stewart [8]), a factorization

$$A(\sigma) = Q_1B(\sigma)Q_2 \tag{21b}$$

with orthogonal Q_1, Q_2 and Hessenberg (or tridiagonal) $B(\sigma)$ is already available, and it may be more economical to factor $B(\sigma)$ instead:

$$B(\sigma) = S'R' \tag{21c}$$

and solve

$$R'\bar{x}^{(0)} = (1, \cdots, 1)^*, \qquad x^{(0)} = Q_2^*\bar{x}^{(0)}, \qquad x^{(0)} = \bar{x}^{(0)}/e^*\bar{x}^{(0)} \tag{5b}$$

in place of (5).

It is a useful fact that the factorization can be reused to find the vector $e^TA(\sigma)^{-1}$ required in (6b) as

$$e^*A(\sigma)^{-1} = e^*R^{-1}S^{-1} \quad \text{using (21a), or} \tag{22a}$$

$$e^*A(\sigma)^{-1} = e^*Q_2^*R'^{-1}S'^{-1}Q_1^* \quad \text{using (21b, c)} \tag{22b}$$

0.1%, respectively, all eigenvalues were found very accurately. We give details for the eigenvalue $\hat{\lambda} = 1$. The associated normalized eigenvector is

$$\hat{x} = (-\tfrac{1}{3640}, 0, \tfrac{1}{364}, 0, -\tfrac{1}{56}, 0, \tfrac{1}{8}, 0, -\tfrac{1}{2}, 0, 1)^*.$$

With constant shift $\sigma = 1.0001$ the sixth iterate $x = x^{(6)}$ was accurate to 15 decimals, e.g.

$$\lambda = 1.000\,000\,000\,000\,000\,1,$$

$$x_{11} = 1 \text{ (normalized)},$$

$$x_{10} = 0.000\,000\,000\,000\,000\,2,$$

$$x_1 = -0.000\,260\,416\,666\,666\,7.$$

This confirms the claim that a single precision factorization coupled with double precision residuals suffices to produce results comparable with the use of double precision throughout. We remark that if (6a) is used in place of (6b) to compute λ_{l+1} then the iteration fails to converge to the small eigenvalues since the corresponding left and right eigenvectors are almost orthogonal.

2. Our second example is a symmetric, definite quadratic eigenvalue problem taken from Scott and Ward [7]:

$$A(\lambda) = \begin{pmatrix} -10\lambda^2+\lambda+10 & & & & \\ 2\lambda^2+2\lambda+2 & -11\lambda^2+\lambda+9 & & \text{sym.} & \\ -\lambda^2+\lambda-1 & 2\lambda^2+2\lambda+3 & -12\lambda^2+10 & & \\ \lambda^2+2\lambda+2 & -2\lambda^2+\lambda-1 & -\lambda^2-2\lambda+2 & -10\lambda^2+2\lambda+12 & \\ 3\lambda^2+\lambda-2 & -\lambda^2+3\lambda-2 & \lambda^2-2\lambda-1 & 2\lambda^2+3\lambda+1 & -11\lambda^2+3\lambda+10 \end{pmatrix}.$$

Its eigenvalues (to three decimals) are:

$$-1.27 \quad -1.08 \quad -1.0048 \quad -.779 \quad -.512$$

$$.502, \quad .880 \quad .937 \quad 1.47 \quad 1.96.$$

Selected results are given in Table 1; listed are

σ—the (constant) shift,
l—the number of iterations (max. 20*),
Δx—max. norm of final eigenvector correction,
q—average quotient of consecutive corrections,
$q^* = \inf_{\lambda \neq \hat{\lambda}} |(\sigma - \hat{\lambda})/(\sigma - \lambda)|$,
$\hat{\lambda}$—the computed eigenvalue.

TABLE 1

σ	l	Δx	$1/q$	$1/q^*$	$\hat{\lambda}$
−1	14	$8_{10^{-11}}$	16.5	15.9	−1.004 838 220 309 025
0	20*	$4_{10^{-7}}$	2.15	1.52	−.511 761 939 586 031 0
.5	8	$2_{10^{-16}}$	208	157	.502 415 273 308 102 5
.9	20*	$7_{10^{-7}}$	1.88	1.82	.879 927 281 097 871 3
.94	16	$6_{10^{-10}}$	16.9	17.4	.936 550 668 659 857 1

It is seen that as described above the convergence rate q strongly correlates with the relative distance q^* of the shift from the eigenvalues; in particular this explains the slow convergence when the shift is near the average of two consecutive eigenvalues.

and to find the correction $dx^{(l)}$ in (8) as

(8a) $$dx^{(l)} = R^{-1}S^{-1}r^{(l)} \quad \text{using (21a), or}$$

(8b) $$dx^{(l)} = Q_2^* R'^{-1} S'^{-1} Q_1^* r^{(l)} \quad \text{using (21b, c).}$$

Equations (8) and (9) suggest that the limit accuracy with which $\hat{x}, \hat{\lambda}$ can be approximated by $x^{(l)}, \lambda^{(l)}$ is mainly determined by the accuracy with which the residual $A(\lambda_{l+1})x^{(l)}$ is computed. Therefore it is sensible to store $x^{(l)}$ and $\lambda^{(l)}$ in double precision. Then λ_{l+1} and $r^{(l)}$ should be computed in double precision, but $r^{(l)}$ can be rounded to single precision before it is stored. The factorizations (21a–c) and the solution of the equations (5a, b), (8a, b) can be performed in single precision, as well as the computation of $e^*A(\sigma)^{-1}$ by (11a, b). Finally, the correction (9) should be done in double precision again. The resulting limit accuracy can then be expected to be about the same as with the use of double precision throughout, and this is confirmed by numerical examples shown below.

Finally, the equations (6a) resp. (6b) need not be solved to full accuracy, and it is sufficient to take for λ_{l+1} one Newton step (linear interpolation) or Euler step (quadratic interpolation) from λ_l (starting with $\lambda_0 := \sigma$) towards the solution.

To demonstrate the behaviour of residual inverse iteration we report here some of the numerical experiments which we have done on the UNIVAC 1100/82 of the University of Freiburg (mantissa length: 27 bits for single precision, 60 bits for double precision). The linear equations were solved using single precision Gauss elimination with column pivoting, and the vector e was chosen as the unit vector with a 1 in the position of the absolutely largest entry of the most recent $x^{(l)}$. This position was found to be independent of l except sometimes for $l = 1$ or 2.

For a fixed shift σ we generally observed global, monotonic and linear convergence of λ_l to one of the eigenvalues nearest to σ. In almost all examples tried, the observed convergence factor for the eigenvector was

$$\frac{\|x^{(l+1)} - \hat{x}\|_\infty}{\|x^{(l)} - \hat{x}\|_\infty} \approx C \cdot \inf_{\lambda \neq \hat{\lambda}} \left| \frac{\sigma - \hat{\lambda}}{\sigma - \lambda} \right|,$$

where the infimum extends over all eigenvalues of $A(\lambda)$ distinct from $\hat{\lambda}$, and C varied between 0.5 and 3. The classical analysis of inverse iteration guarantees such a behaviour in the linear, nondefective case.

Although our convergence analysis applies only to simple, isolated eigenvalues it was found that multiple, nondefective eigenvalues were found with the same speed and accuracy as simple eigenvalues. We did not try residual inverse iteration on defective problems. We now consider specific examples.

1. Our first example is a standard eigenvalue problem $A_0\hat{x} = \hat{\lambda}\hat{x}$, corresponding to $A(\lambda) = A_0 - \lambda I$. The matrix A_0 is the Frank matrix of order 11 (see [3]):

$$A_0 = \begin{pmatrix} 11 & 10 & 9 & \cdots & 1 \\ 10 & 10 & 9 & \cdots & 1 \\ & 9 & 9 & \cdots & 1 \\ 0 & & \ddots & \ddots & \vdots \\ & & & 1 & 1 \end{pmatrix}.$$

All eigenvalues are simple. With strategy (6b) and constant shifts accurate to 10% and

3. Our third example is another symmetric definite problem of Scott and Ward [7], this time with multiple eigenvalues:

$$A(\lambda)=\begin{pmatrix} -\lambda^2-3\lambda+1 & & & \text{sym.} \\ \lambda^2-1 & -2\lambda^2-3\lambda+5 & & \\ -\lambda^2-3\lambda+1 & \lambda^2-1 & -2\lambda^2-5\lambda+2 & \\ -2\lambda^2-6\lambda+2 & 2\lambda^2-2 & -4\lambda & -9\lambda^2-19\lambda+14 \end{pmatrix}.$$

This matrix function has the double eigenvalues 1 and -2, and the simple eigenvalues

$$-4\pm\sqrt{19}\ (\approx -8.36, .359), \qquad -4\pm\sqrt{18}\ (\approx -8.24, .243).$$

Sample results are given in Table 2.

TABLE 2

σ	l	Δx	$1/q$	$1/q^*$	$\hat{\lambda}$
.3	20*		1.05	1.03	.247 ···
.25	14	$2_{10^{-17}}$	15.3	14.8	.242 640 687 119 285 1
1.01	9	$5_{10^{-17}}$	57	65.1	1
−2.01	7	$7_{10^{-17}}$	413	623	2

The multiple eigenvalues are found as efficiently as the others.

4. Finally, we demonstrate the effect of convergence acceleration with the matrix function of Example 2. Starting with $\sigma=2$, the shift was updated in each iteration, replacing σ by the single precision truncation of the most recent λ_l. After 8 steps the approximate eigenpair agrees to 16 decimal places with that computed by the algorithm with constant shift $\sigma=1$; but the λ_l are no longer monotonic, and the limit eigenvalue $-1.0048\cdots$ is no longer nearest to the initial shift. The convergence behaviour can be seen from Table 3 which lists the maximal element of the residuals and the eigenvector corrections.

TABLE 3

Step	Residual	Correction
1	4.77	$1.85_{10^{-1}}$
2	4.63	$1.13_{10^{0}}$
3	1.07	$3.37_{10^{-1}}$
4	$2.14_{10^{-2}}$	$9.77_{10^{-3}}$
5	$6.82_{10^{-7}}$	$3.70_{10^{-4}}$
6	$1.12_{10^{-9}}$	$7.75_{10^{-10}}$
7	$1.19_{10^{-16}}$	$6.64_{10^{-17}}$
8	$5.51_{10^{-17}}$	$5.64_{10^{-18}}$

REFERENCES

[1] J. J. DONGARRA, C. B. MOLER AND J. H. WILKINSON, *Improving the accuracy of computed eigenvalues and eigenvectors*, this Journal, 20 (1983), pp. 23–45.

[2] D. K. FADDEEV AND V. N. FADDEEVA, *Computational Methods of Linear Algebra*, Freeman, San Francisco, 1963.

[3] R. T. GREGORY AND D. L. KARNEY, *A Collection of Matrices for Testing Computational Algorithms*, Wiley-Interscience, New York–London, 1969.

[4] W. GRÖBNER, *Matrizenrechnung*, Bibliogr. Inst., Mannheim - Wien - Zürich, 1966.

[5] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, Prentice-Hall, Englewood Cliffs, NJ, 1980.

[6] A. RUHE, *Algorithms for the nonlinear eigenvalue problem*, this Journal, 10 (1973), pp. 674–689.

[7] D. S. SCOTT AND R. C. WARD, *Solving symmetric-definite quadratic problems without factorization*, SIAM J. Sci. Stat. Comput., 3 (1982), pp. 58–67.

[8] G. W. STEWART, *Introduction to Matrix Computations*, Academic Press, New York - San Francisco - London, 1973.

[9] H. J. SYMM AND J. H. WILKINSON, *Realistic error bounds for a simple eigenvalue and its associated eigenvector*, Numer. Math., 35 (1980), pp. 113–126.

[10] H. UNGER, *Nichtlineare Behandlung von Eigenwertaufgaben*, Z. Angew. Math. Mech., 30 (1950), pp. 281–282.

[11] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford Univ. Press, London, 1965.

[12] T. YAMAMOTO, *Error bounds for computed eigenvalues and eigenvectors*, Numer. Math., 34 (1980), pp. 189–199.