Refining the General Symmetric Definite Eigenproblem

1-Page Overview:

$A = A'$, $H = H'$ is +ve def.; $A \cdot E = H \cdot E \cdot \Lambda$ for diag. Λ of eig'ues.

Direct Method: Cholesky factorize $H = U' \cdot U$; get $W := U'^{-1} \cdot A \cdot U^{-1} = W'$

Eigenvalues $\Lambda := Q' \cdot W \cdot Q$ for $Q = Q^{-1}$; Eigenvectors $E := U^{-1} \cdot Q$

Iteration: $A_k := (E_1 \cdot E_2 \cdot \ldots \cdot E_k)' \cdot A \cdot (E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to \Lambda$ as $k \to \infty$;

$H_k := (E_1 \cdot E_2 \cdot \ldots \cdot E_k)' \cdot H \cdot (E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to \mathrm{Diag}(H_k) := I$; !

each $E_k - I$ is "2-by-2" to zero off-diag. pairs; $(E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to E$.

Convergence: Long known to be ***Locally Quadratic*** .
Here proved ***Global*** . … via a monotonic determinant !

Numerical STABILITY THREATENED by *PATHOLOGIES*:

- Near-repeated roots ***Ruin*** E_k computed in the obvious way.

  Unobvious formulas for E_k fend off this and another pathology:
- Nearly SINGULAR H causes some HUGE eigenvector(s).
- When A and H share a *Near-Nullspace*, some computed eigenvectors can be Huge Unnecessarily, thus exacerbating Roundoff's Corruption of eigenvalues/vectors otherwise Well-Determined by the data. … remedied by an iterative refinement

Question: How likely would we notice **Remediable Corruption**?

……

Test Data include N-by-N segments $H_{N,K}$ of ***Hilbert Matrices***:

$$\{ H_{N,K} \}_{i,j} = 1/(i+j+K-1) .$$

We provide preponderantly *Integer* (exact!) formulas for the four Cholesky factors of $H_{N,K}$ and $H_{N,K}^{-1}$, and for the *Bidiagonal* whose (singular values)$^2 = \Lambda$ when $A := H_{N,K+1}$ & $H := H_{N,K}$.

……

This document is posted at <www.eecs.berkeley.edu/~wkahan/HHXVIII.pdf>. For details see <…/Math128/GnSymEig.pdf> and <…/MathH110/HilbMats.pdf> .

Iteration: $A_k := (E_1\cdot E_2\cdot\ldots\cdot E_k)'\cdot A\cdot(E_1\cdot E_2\cdot\ldots\cdot E_k) \to \Lambda$ as $k \to \infty$;
$H_k := (E_1\cdot E_2\cdot\ldots\cdot E_k)'\cdot H\cdot(E_1\cdot E_2\cdot\ldots\cdot E_k) \to \mathrm{Diag}(H_k) := I$; **!**
each $E_k - I$ is "2-by-2" to zero off-diag. pairs; $(E_1\cdot E_2\cdot\ldots\cdot E_k) \to E$.

How will each E_k be computed?

Squash E_k to a 2-by-2 E; want diagonals $E'\cdot A\cdot E = \Lambda$ & $E'\cdot H\cdot E = I$

given squashed submatrices $A := \begin{bmatrix} v_1 & \alpha \\ \alpha & v_2 \end{bmatrix}$ & $H := \begin{bmatrix} 1 & \sigma \\ \sigma & 1 \end{bmatrix}$ with $-1 < \sigma < 1$.

Obvious Method: $\Lambda = \mathrm{Diag}([\lambda_1, \lambda_2])$ needs zeros λ_j of ***Char. Poly.***:

$$\mathrm{Det}(\lambda\cdot H - A) = (1-\sigma^2)\cdot\lambda^2 - (v_1 + v_2 - 2\sigma\cdot\alpha)\cdot\lambda + v_1\cdot v_2 - \alpha^2 .$$

Discriminant: $\delta^2 := (1-\sigma^2)\cdot(v_1 - v_2)^2 + \big(\sigma\cdot(v_1 + v_2) - 2\alpha\big)^2 \ge 0$.

Zeros $\lambda_j := \frac{1}{2}(v_1 + v_2 - 2\sigma\cdot\alpha \pm \delta)/(1-\sigma^2)$ ordered so $\dfrac{\lambda_1 - \lambda_2}{v_1 - v_2} \ge \dfrac{1}{\sqrt{1-\sigma^2}}$.

Then $E := \begin{bmatrix} \sqrt{\lambda_1 - v_2} & \dfrac{\sigma\cdot\lambda_2 - \alpha}{\sqrt{v_1 - \lambda_2}} \\ \dfrac{\alpha - \sigma\cdot\lambda_1}{\sqrt{\lambda_1 - v_2}} & \sqrt{v_1 - \lambda_2} \end{bmatrix} / \sqrt{(|\delta|\cdot\mathrm{sgn}(\lambda_1 - \lambda_2))}$. (sgn = ±1)

2 Pathologies Exacerbate Roundoff to Ruin E_k and All that Follows:

- Near-Repeated Zeros λ_j: $\lambda_1 \approx v_1 \approx v_2 \approx \lambda_2$, so $E \approx O/0$.
- Near-Singular H: $|\sigma| \approx 1$, so at least one λ_j is *Huge*; $E \approx \infty$.

Numerical examples in GnSymEig.pdf's §3 show that the threat is real.

(How might suspicions about the accuracy of a computed E be allayed or else corroborated by someone who cannot perform an error-analysis? GnSymEig.pdf's §3 illustrates how my …/Mindless.pdf's §14 does it.)

How will each E_k be computed *Accurately* ?

Squash E_k to a 2-by-2 E ; want diagonals $E'\cdot A\cdot E = \Lambda$ & $E'\cdot H\cdot E = I$

given squashed submatrices $A := \begin{bmatrix} v_1 & \alpha \\ \alpha & v_2 \end{bmatrix}$ & $H := \begin{bmatrix} 1 & \sigma \\ \sigma & 1 \end{bmatrix}$ with $-1 < \sigma < 1$.

Unobvious Method: … in GnSymEig.pdf's §2

Case $|\sigma| \le 3/4$, when 2-by-2 H is not nearly singular:

First compute $\Theta := \frac{1}{2}\arcsin(\sigma)$ and $\Phi := \frac{1}{2}\arctan\left(\frac{2\alpha - (v_1 + v_2)\cdot\sigma}{(v_1 - v_2)\cdot\cos(2\Theta)}\right)$,

choosing $|\Phi| \le \pi/4$ and letting $0/0 := 0$. Then

$$E := \begin{bmatrix} \cos(\Phi + \Theta) & -\sin(\Phi + \Theta) \\ \sin(\Phi - \Theta) & \cos(\Phi - \Theta) \end{bmatrix} / \cos(2\Theta) \ .$$

(E can be computed from strictly algebraic operations $\{+, -, \cdot, /, \sqrt{}\}$,
but perhaps slightly less accurately than from a modern Math library.)

Case $3/4 \le |\sigma| \le 1$, when 2-by-2 H may be nearly singular:
This case tends to arise in early iterations deserving exact cancellations of data.

First set $V := \begin{bmatrix} \sqrt{2 + 2|\sigma|} & 0 \\ 0 & \sqrt{2 - 2|\sigma|} \end{bmatrix}$ and $X := \begin{bmatrix} 1 & \mathrm{sgn}(\sigma) \\ \mathrm{sgn}(\sigma) & -1 \end{bmatrix}$. (sgn = ±1)

Next $T := V^{-1}\cdot(X\cdot(A\cdot X))\cdot V^{-1}$; $\tau := \frac{1}{2}\arctan\left(2\cdot t_{12}/(t_{11} - t_{22}) \right)$;

$\psi := \tau + (\mathrm{sgn}(\sigma)\cdot\mathrm{sgn}(\tau) - 1)\cdot\mathrm{sgn}(\tau)\cdot\pi/4$; $Y := \begin{bmatrix} \cos(\psi) & \sin(\psi) \\ \sin(\psi) & -\cos(\psi) \end{bmatrix}$;

$E := X\cdot V^{-1}\cdot Y$ and $\Lambda := Y\cdot T\cdot Y = \mathrm{Diag}([\lambda_1, \lambda_2])$ ordered as before.

If τ is indeterminate because of … arctan(0/0) … just set $\psi := 0$.

Numerical examples and tests in GnSymEig.pdf's §3 corroborate the numerical stability of this page's E and Λ versus the previous page's.
And all formulas' $\mathrm{Det}(E) = 1/\sqrt{(1 - \sigma^2)} \ge 1$.

Iteration: $A_k := (E_1 \cdot E_2 \cdot \ldots \cdot E_k)' \cdot A \cdot (E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to \Lambda$ as $k \to \infty$;
$H_k := (E_1 \cdot E_2 \cdot \ldots \cdot E_k)' \cdot H \cdot (E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to \mathrm{Diag}(H_k) := I$; **!**
each $E_k - I$ is "2-by-2" to zero off-diag. pairs; $(E_1 \cdot E_2 \cdot \ldots \cdot E_k) \to E$.

How do we know that this iteration converges?

Every $\mathrm{Det}(E_k) \geq 1$. Therefore every $\mathrm{Det}(H_k) \leq \mathrm{Det}(H_{k+1}) \leq 1$ because we keep every $\mathrm{Diag}(H_k) := I$, and so $\mathrm{Det}(H_k) \to$ (a positive limit) ≤ 1 .

Lemma: If h is the biggest magnitude of the off-diagonal elements of a positive definite N-by-N matrix H whose diagonal $\mathrm{Diag}(H) = I$, then $0 < \det(H) \leq 1 - h^2$. If also $0 < h \cdot (N-1) < 1$, $\det(H) > \left(1 - (N-1) \cdot h\right)^N$.

So long as the iteration visits every off-diagonal element infinitely often and annihilates just those not too much smaller than average at the time, a compactness argument goes from "$\mathrm{Det}(H_k)$ converges" to "$\mathrm{Det}(E_k) \to 1$" to "$\mathrm{Det}(H_k) \to 1$" to "$H_k \to I$" to "$A_k \to \Lambda$" .

For details see §6 of GnSymEig.pdf; its §11 offers an experimental MATLAB program that implements a strategy for choosing which pairs of off-diagonal elements to annihilate.

When should iteration stop? Hard to know for sure. Small residuals can mislead. Suppose $\mathbf{v}$ approximates an eigenvector. Then the *Rayleigh Quotient* $\omega := \mathbf{v}' \cdot A \cdot \mathbf{v} / \mathbf{v}' \cdot H \cdot \mathbf{v}$ approximates an eigenvalue better. How well? The residual $\mathbf{r} := A \cdot \mathbf{v} - \omega \cdot H \cdot \mathbf{v}$ figures in an eigenvalue estimate $\omega \pm \sqrt{(\mathbf{r}' \cdot H^{-1} \cdot \mathbf{r} / \mathbf{v}' \cdot H \cdot \mathbf{v})}$ that is costly to compute and can be wide, though $\mathbf{r}$ is tiny, when H is too nearly singular.

Our Chosen Normalizations

- Diag_of(H) = I , achieved by an N-by-N diagonal congruence.
  Why? Among all diagonal congruences, the unknown one that minimizes the condition number $||H||\cdot||H^{-1}||$ brings it down no smaller than 1/N of the condition number of what we have chosen (A. van der Sluis, 1969).

- Eigenvector matrix E has $E'\cdot H\cdot E = I$ when $E'\cdot A\cdot E = \Lambda$ is diagonal, achieved by scaling the eigenvectors (columns of invertible E).
  Why? Convenience. Then Λ is a diagonal of eigenvalues. Moreover then $||E|| = \sqrt{||H^{-1}||}$, so every eigenvector **e** has $||\mathbf{e}|| \le ||H^{-1}||$ and at least one has $||\mathbf{e}|| \ge \sqrt{(||H^{-1}||/N)}$, thus exhibiting roughly how close H is to singular.

Sensitivity to Infinitesimal Perturbations ∂…

Perturbing $A \to A + \partial A$ and $H \to H + \partial H$ changes a simple eigenvalue $\lambda \to \lambda + \partial\lambda$ and its eigenvector $\mathbf{e} \to \mathbf{e} + \partial\mathbf{e}$. Besides $A\cdot\mathbf{e} = \lambda\cdot H\cdot\mathbf{e}$, our chosen normalization has $\mathbf{e}'\cdot H\cdot\mathbf{e} = 1$ so $\lambda = \mathbf{e}'\cdot A\cdot\mathbf{e}$. Differentiate to get

$$\partial\lambda = \mathbf{e}'\cdot(\partial A - \lambda\cdot\partial H)\cdot\mathbf{e} .$$

This implies that λ is hypersensitive to tiny perturbations *only* when $||\mathbf{e}||$ is big even if $|\lambda|$ is small. Apparently this hypersensitivity extends also to perturbations due to roundoff (though "*only*" must be omitted for some numerical methods, like the obvious formula for E_k on p. 2).

$||\mathbf{e}||$ can be big only if H is nearly singular, and then this **e**'s eigenvalue λ is deservedly "ill-conditioned" (hypersensitive to tiny perturbations and roundoff), and especially so if λ is among the smaller eigenvalues.

When can an ill-conditioned $|\lambda| = |\mathbf{e}'\cdot A\cdot\mathbf{e}|$ be small though $||\mathbf{e}||$ is big?
- When $||A||$ is small enough so are all eigenvalues, including this λ .
- When A is indefinite, $\lambda = \mathbf{e}'\cdot(A\cdot\mathbf{e})$ can cancel though $||A\cdot\mathbf{e}||$ is big.
- When A and H share a *Near-Nullspace*, λ may be small, or not.

This last case sometimes causes (no matter how) *computed* eigenvalues and eigenvectors to appear ill-conditioned undeservedly. Why? …

When A and H Share a *Near-Nullspace* …

Then some $\mathbf{z} \neq \mathbf{o}$ satisfies $||A\cdot\mathbf{z}||/||\mathbf{z}|| << ||A||$ and $||H\cdot\mathbf{z}||/||\mathbf{z}|| << ||H||$.

Then some multiple $\zeta\cdot\mathbf{z}$ can be added to a computed approximation $\mathbf{f}$ to an eigenvector, satisfying $A\cdot\mathbf{f} \approx \omega\cdot H\mathbf{f}$ and $\mathbf{f}'\cdot H\mathbf{f} = 1$ so $\omega \approx \mathbf{f}'\cdot A\mathbf{f}$ estimates $\mathbf{f}$'s eigenvalue, leaving these equations roughly equally well satisfied by $\mathbf{f} + \zeta\cdot\mathbf{z}$ as by $\mathbf{f}$ though these may differ substantially.

Which of $\mathbf{f}$ and $\mathbf{f} + \zeta\cdot\mathbf{z}$ is the better approximation to an eigenvector?

If $\mathbf{f}$ is so contaminated by $\mathbf{z}$ that $||\mathbf{f}||$ is much bigger than necessary, then so is the estimated sensitivity $\partial\omega = \mathbf{f}'\cdot(\partial A - \omega\cdot\partial H)\cdot\mathbf{f}$; and in some numerical experiments roundoff spoiled ω much more than necessary.

When $\mathbf{f}$ was replaced by $\boldsymbol{f}$ obtained from $\mathbf{f} - \mathbf{z}\cdot(\mathbf{z}'\cdot H\mathbf{f})/(\mathbf{z}'\cdot H\mathbf{z})$ after scaling it to make $\boldsymbol{f}'\cdot H\boldsymbol{f} = 1$, the new estimate $\varpi := \boldsymbol{f}'\cdot A\boldsymbol{f}$ was more accurate than ω except when cancellation left $||\mathbf{f} - \mathbf{z}\cdot(\mathbf{z}'\cdot H\mathbf{f})/(\mathbf{z}'\cdot H\mathbf{z})||$ too much smaller than $||\mathbf{f}||$, which case called for iterative refinement.

But how do we tell which is $\mathbf{f}$ and which is $\mathbf{z}$?

More generally, suppose the computed eigenvector matrix intended to approximate E is F satisfying $F'\cdot H\cdot F \approx I$ and $F'\cdot A\cdot F \approx W$ is diagonal. If H is nearly singular, at least one column of F must be big. If more than one column is big, group the big columns into a submatrix B and obtain its Economical Singular Value Decomposition $B = P\cdot\Psi\cdot Q$ in which $P'\cdot P = I$, the diagonal Ψ of singular values, and $Q' = Q^{-1}$ all have the same small (we hope) dimensions. Typically $B\cdot Q' = P\cdot\Psi$ will have one or two big columns and the rest small. Rescale its columns to get $\overline{B}$ satisfying $\mathrm{Diag\_of}(\overline{B}'\cdot H\cdot\overline{B}) = I$. Then apply the iteration of pp. 3-4 to diagonalize $\overline{B}'\cdot H\cdot\overline{B}$ and $\overline{B}'\cdot A\cdot\overline{B}$ by simultaneous congruence. This is the iterative refinement. It works better if $\overline{B}'\cdot H\cdot\overline{B}$ and $\overline{B}'\cdot A\cdot\overline{B}$ are computed extra-precisely.

Does that process always work? That process is a WORK IN PROGRESS.

Four Cholesky Factors of Hilbert Matrices and their Inverses

Formulas for them, derived long ago, were resurrected to use in test data for investigations into the General Symmetric Definite Eigenproblem. The formulas and MATLAB programs based upon them are posted at
<www.eecs.berkeley.edu/~wkahan/MathH110/HilbMats.pdf> .

N-by-N Hilbert matrices $H_{N,K}$ have elements $\{H_{N,K}\}_{ij} = 1/(i+j+K-1)$ for integers $K \geq 0$. When the quotients' rounding errors are intolerable we use instead the integer matrix $Y_{N,K} = L_{N,K}\cdot H_{N,K}$ computed exactly unless integer $L_{N,K} :=$ LCM([K+1, K+2, …, 2N+K–1]) gets rounded off.

Integer matrix $H_{N,K}^{-1}$ is computed from recurrences first published by Dr. Sam Schechter in *MTAC* in 1959. I cannot find my old notes with derivations and proofs of the other formulas. They use an old notation:

Combinatorial/Binomial Coefficient $^{n}Ç_{k} := n!/(\,k!\cdot(n-k)!\,)$

N-row $\mathbf{u}'$:= [1, 1, 1, …,1, 1]

N-by-N Diagonal \$:= Diag([1, -1, 1, -1, …, $(-1)^{N-1}$])

Henceforth subscripts N,K are taken for granted so that the abbreviation H can be used for $H_{N,K}$ and likewise for all matrices except $\mathbf{u}'$ and \$.

Diagonal D has $\{D\}_{j,j} := d_{N,K,j} := (-1)^j\cdot j\cdot {}^{N}Ç_{j}\cdot {}^{N+K+j-1}Ç_{N}$.

$H^{-1} = D\cdot H\cdot D$ has integer elements $\{H^{-1}\}_{i,j} = d_{N,K,i}\cdot d_{N,K,j}/(i+j+K-1)$

$\det(H^{-1}) = |\det(D)| = |\prod_j d_{N,K,j}|$

$\mathbf{u}'\cdot H^{-1}\cdot\mathbf{u} := \sum_i\cdot\sum_j \{H^{-1}\}_{i,j} = N\cdot(N+K)$ is used to test H^{-1}'s accuracy.

Approximations valid as $K \to \infty$:

$$\mathbf{u}'\cdot\$\cdot H^{-1}\cdot\$\mathbf{u} := \sum_i\cdot\sum_j |\{H^{-1}\}_{i,j}| = \left(\,4^{N-1}\cdot(N+K)^{2N-1}/((N-1)!)^2\,\right)\cdot\left(1 + O((N+K)^{-2})\,\right)$$

$$\|H^{-1}\|_F^2 := \sum_i\cdot\sum_j (\{H^{-1}\}_{i,j})^2 = \left(\,(2N-2)!\cdot(N+K)^{2N-1}/((N-1)!)^4\,\right)^2\cdot\left(1 + O((N+K)^{-2})\,\right)$$

Four N-by-N upper-triangles U and R and their inverses are Cholesky factors of $H = U'\cdot U = R^{-1}\cdot R'^{-1}$ and of $H^{-1} = R'\cdot R = U^{-1}\cdot U'^{-1}$. Though not generally integer matrices, they can be assembled out of integer matrices starting with $ and three more N-by-N diagonals:

$$\{¥\}_{j,j} := K + 2j - 1 ; \qquad \{£\}_{j,j} := {}^{K+2j-2}Ç_{j-1} ; \qquad \{Ω\}_{j,j} := {}^{K+N-1+j}Ç_{N-j} .$$

These combine with two N-by-N integer upper triangles C and G :

$$\{C\}_{i,j} := {}^{2j-1+K}Ç_{j-i} ; \quad G := \$\cdot¥\cdot C^{-1}\cdot¥^{-1}\cdot\$ \quad \text{so} \quad \{G\}_{i,j} := {}^{i+j-2+K}Ç_{j-i} .$$

Their combinations produce the four Cholesky factors:

$$U = \sqrt{¥}\cdot C\cdot(¥\cdot£)^1 ; \qquad U^{-1} = £\cdot\$\cdot G\cdot\$\sqrt{¥} ;$$
$$R = \$\cdot\sqrt{¥}\cdot C\cdot\$Ω = \$\cdot U\cdot D ; \quad R^{-1} = (Ω\cdot¥)^{-1}\cdot G\cdot\sqrt{¥} .$$

These formulas seem at first to demand $O(N^3)$ arithmetic operations, but the MATLAB programs in HilbMats.pdf use recurrences that need only $O(N^2)$ operations.

HilbMats.pdf also offers tests of the MATLAB programs' correctness. The tests would be simpler and more reliable if MATLAB supported IEEE Standard 754's exception flags, particularly the Inexact Flag.

……

When $A := H_{N,K+1}$ and $H := H_{N,K}$ in the General Symmetric Definite Eigenproblem $A\cdot E = H\cdot EΛ$, the eigenvalues in Λ are the squared singular values of the upper-triangular $F_{N,K} := U_{N,K+1}\cdot U_{N,K}^{-1}$, which turns out to be bidiagonal and computable very quickly from the N-by-N bidiagonal J_N whose elements $\{J_N\}_{i,j} :=$ (if $i \le j \le i+1$ then i else 0) .

$$F_{N,K} = (\sqrt{¥}_{N,K+1})^{-1}\cdot(J_N + K\cdot I)\cdot(\sqrt{¥}_{N,K})^{-1} .$$

And its singular values are computable quickly and very accurately.

……